Journal of Machine Learning Research 1 (2013) 1-60 Submitted 12/13; Published 00/00

Predictive Interval Models for Non-parametric Regression

Mohammad Ghasemi Hamed GHASEMI@RECHERCHE.ENAC.FR
ENAC, MAIAA, F-31055 Toulouse, France

Mathieu Serrurier SERRURIER@IRIT.FR
IRIT - Université Paul Sabatier
Univ. de Toulouse, IRIT, F-31400 Toulouse, France

Nicolas Durand DURAND@RECHERCHE.ENAC.FR
IRIT - Université Paul Sabatier
ENAC, MAIAA, F-31055 Toulouse, France

Editor:

arXiv:1402.5874v1 [cs.LG] 24 Feb 2014

Abstract

Having a regression model, we are interested in finding two-sided intervals that are guaranteed to contain at least a desired proportion of the conditional distribution of the response variable given a specific combination of predictors. We name such intervals predictive intervals. This work presents a new method to find two-sided predictive intervals for non-parametric least squares regression without the homoscedasticity assumption. Our predictive intervals are built by using tolerance intervals on prediction errors in the query point's neighborhood. We proposed a predictive interval model test and we also used it as a constraint in our hyper-parameter tuning algorithm. This gives an algorithm that finds the smallest reliable predictive intervals for a given dataset. We also introduce a measure for comparing different interval prediction methods yielding intervals having different size and coverage. These experiments show that our methods are more reliable, effective and precise than other interval prediction methods.

Keywords: Interval prediction, tolerance intervals, predictive intervals, local linear regression

1. Introduction

Regression analysis is a statistical technique for estimating the value of one variable as a function of independent variables. They are widely applied in science and engineering, they are used in problems like function estimation, financial forecasting, and time series prediction. In the general form a regression equation has three variables: the response variable $Y(x)$, a function $f(x)$ and a random error ε, where $Y(x) = f(x) + \varepsilon$. There are different kinds of regression techniques which estimate different characteristics of the conditional distribution of the response variable $Y(x)$. The most common approaches estimate the mean of the random variable $Y(x)$ and are usually known as least-squares techniques. Robust regression approaches are similar to least-squares techniques but they are designed to be

©2013 Mohammad Ghasemi Hamed and Mathieu Serrurier and Nicolas Durand .

robust to outliers and violations of the least-squares assumptions. Another kind, called quantile regression, estimates the conditional quantiles of the response variable. In each category, the regression function can be estimated with a parametric linear, a parametric non-linear or a non-parametric method. This results in linear, non-linear or non-parametric regression models.

Interval prediction is an important part of every regression modeling procedure because regression models are always built with finite sample sizes. Thus the predicted mean or quantile is an estimate of the true unknown conditional mean or quantile of the random variable $Y(x)$. Therefore while dealing with finite size datasets, we need to make some statistical inferences. There are currently many interval prediction methods for the regression context, however each interval has its own specific interpretation and application. In this work we are interested in finding two-sided prediction intervals in regression models which contain, with a high confidence level, at least a desired proportion of the conditional response variable. Such intervals are mostly used in application demanding a high level of confidence, like quality control, environmental monitoring, industrial hygiene, exposure data analysis, aircraft trajectory prediction, security systems etc.

We divide interval prediction approaches in regression into two main categories: The first category methods are based on the estimated conditional mean. These methods are usually based on least-squares models and propose interval prediction techniques that are centered on the estimation of the mean regression function. These approaches generally assume a non-biased regression model with a Gaussian error having constant variance. On the other hand we have quantile regression methods which directly estimate these intervals. Quantile regression methods are more robust to outliers and have less assumptions than the least-squares approaches. But they suffer from other weaknesses like slower speed of convergence and the crossing quantile effect (two different quantiles may cross or overlap each other).

One of our contributions is the review and the comparison of different least-squares and quantile regression techniques used to find intervals which contain a desired proportion of the response variable. We take advantage of this work to address common misunderstood questions about interval prediction methods in the machine learning community. We explain their applications and review their drawbacks. As pointed out at the beginning paragraph, we are interested in finding intervals in regression models which contain, with a high confidence level, at least a desired proportion of the conditional response variable. For that purpose, we introduce a new type of interval prediction method named "predictive interval methods". A predictive interval model is guaranteed to contain for any query point x, at least a desired proportion of the conditional distribution of the response variable. Such models can be obtained with tolerance intervals for regression or confidence interval on quantile regression, but these concepts have limited applications in the literature. So we propose tolerance intervals for Local Linear Regression (LLR). Then we use the tolerance intervals for LLR to obtain predictive interval models for LLR. Our predictive interval models are applied for two-sided interval prediction, however one can easily extend them to a one-sided interval prediction context. Then, we introduce a statistical test to check if an "interval prediction model" is a "predictive interval model". In the same context, we

introduce two measures for ranking interval prediction models. These measures rate the efficiency and the tightness of the obtained envelope. Our predictive interval methods are based on LLR and give variable size intervals. We assume that the mean regression function is locally linear and that the prediction error is locally homoscedastic (heterocedastic in general). Our method does not neglect the regression bias and finds intervals that work properly with biased regression models. The proposed predictive intervals are based on the Leave-One-Out (LOO) or 10-fold cross-validation prediction errors of the local linear regression.

In order to validate our methods, we use several regression datasets to compare our predictive interval method for local linear regression with other interval prediction methods. The selected methods will be tested on their capacity to provide two-sided β-content predictive interval models. The models are compared by their reliability, efficiency, precision and the tightness of their obtained envelope. We also take advantage of our evaluation chapter to show that the conventional interval prediction method is not appropriate for high confidence interval prediction. It is almost always less efficient than our predictive interval methods and their envelope is almost always larger than the envelope obtained by our methods. Ghasemi Hamed et al. (2012) proposed a K-Nearest Neighbors (KNN) based interval prediction method which finds intervals that guarantee to contain, simultaneously for all instances in the predictor space, at least a proportion β of the conditional distribution of the response value. They called it simultaneous interval regression for KNN and they stated that, their method is similar to simultaneous tolerance intervals for KNN regression with a high confidence level. Their work differs from our contribution in at least three points: first, we do not look for models that guarantee the simultaneous condition. Our models could be obtained by tolerance intervals for regression and not simultaneous tolerance intervals for regression. The difference between these concepts is explained in Section 3. Second, our methods are based on LLR and use local distribution of the prediction error instead of local distribution of the conditional response variable. Third, we propose a test and two interval comparison measures to obtain and compare our models.

This paper is organized as follows: Section 2 is a background on regression and tolerance interval for least squares regression. Section 3 describes the state of the art in for interval prediction in regression. Section 4 introduces two efficiency measures which can be used to compare different interval prediction models. Section 5 introduces the concept of predictive interval models. Section 6 proposes a methods to obtain tolerance intervals for LLR under the Local Homoscedastic Normal Prediction Error (LHNPE) conditions. Section 7 describes how to use the tolerance intervals for LLR to find predictive interval models for LLR. Chapter 8 use experiments to compares our methods with other least squares and quantile regression method on nine benchmark datasets and the final section is a discussion with concluding remarks.

2. Background

2.1 Context and Notation

We choose a fixed regression design where dataset $\mathcal{S} = (x_1, Y(x_1)), \cdots, (x_n, Y(x_n))$ is a random sample. The x_i's are vectors of observations and $Y(x_i)$ are random variables. These distributions are continuous probability distributions. We always suppose that there is one true mean regression function $f(\cdot)$ with a zero mean error and an unknown variance σ^2. The most practical assumption is the Gaussian homoscedastic error, but it is not mandatory. $\mathcal{S}$ is a finite random sample, so the estimated regression model finds a pair of $(\hat{f}, \hat{\sigma})$; $\hat{f}$ denotes the estimated regression function and $\hat{\sigma}$ is the estimated error standard deviation. This pair is a random vector in the probability space of regression models defined for the underlying regression type (for ex: Ordinary Least Squares (OLS)). Note that in the case of error being not normally distributed, the pair $(\hat{f}, \hat{\sigma})$ does not correctly represent the estimated regression model. Thus we will use the symbol $P_{\mathcal{S}}$ instead of $P_{\hat{f},\hat{\sigma}}$ to refer to a probability distribution where the random vector is the estimated regression model based on the random sample $\mathcal{S}$. We also use the following notation:

- $\mathcal{S} = (x_1, Y(x_1)), \cdots, (x_n, Y(x_n))$: the random sample of regression, the training set;
- $f(\cdot)$: the true and unknown regression function;
- $f(x)$: the conditional mean of the response variable for specified combination of the predictors;
- $\hat{f}(\cdot)$: the estimated regression function;
- $\hat{f}(x)$: the estimated regression function at point x;
- ε: the error variable;
- ε_x^{pred}: the prediction error at point x, $\varepsilon_x^{pred} = Y(x) - \hat{f}(x)$;
- σ^2: the true and unknown variance of the error variable;
- $\hat{\sigma}^2$: the estimated variance of the error variable;
- $\sigma^2_{\hat{f}(x)}$: the variance of the estimated regression function at point x;
- x^*: a point in the predictor space that does not exists in the training set;
- $Y(x)$: the conditional response variable for a given combination of the predictors, $Y(x) = f(x) + \varepsilon$;
- Y_i: the i^{th} random response variable, $Y_i = Y(x_i)$;
- y_i: an observation of the random variable Y_i;
- $I(x)^T_{\gamma,\beta}$: β-content γ-coverage tolerance interval for the distribution of the response variable at point x;

- $I(\varepsilon_x^{pred})_{\gamma,\beta}^T$: β-content γ-coverage tolerance interval for the distribution of the prediction error at point x;
- $I(x)_\beta^P$: β-content predictive interval for the distribution of $Y(x)$;
- $I(\varepsilon_x^{pred})_\beta^P$: β-content predictive interval for the distribution of the prediction error at point x;
- Z_p: the p-quantile of a standard normal distribution;
- $\chi^2_{p,n}$: the p-quantile of a chi-square distribution with n degrees of freedom.

  Note that, in this work we suppose that the prediction error for any point x_i, is obtained with $\varepsilon_{x_i}^{pred} = Y(x_i) - \hat{f}(x_i)$, where the mean estimate $\hat{f}(x_i)$ does not use the observation (x_i, y_i).

2.2 Least-squares Regression

Regression analysis is a statistical technique for estimating the value of one variable as a function of independent variables. In fixed design regression, there are n pairs of observations $(x_1, Y_1), \cdots, (x_n, Y_n)$, where x_i is the vector of observations known as covariates and Y_i is the response variable. In other words, the random variable Y_i or $Y(x_i)$ follows a mean function $f(x_i)$ with a random error term ε_i defined as:

$$Y_i = f(x_i) + \varepsilon_i, \text{where } E(\varepsilon_i) = 0. \tag{1}$$

The model assumes that ε_i are mutually independent and identically distributed (iid) random variables. The goal is to estimate the mean function $f(\cdot)$ by $\hat{f}(\cdot)$, being as close as possible to the unknown function $f(\cdot)$. The usual assumption is to suppose that the variance of the error is the same everywhere. This is known as homoscedasticity and the opposite hypothesis (variable error variance) is known as heteroscedasticity.

In a least squares regression, the idea is to estimate the mean of $Y(x)$ by $\hat{f}(x)$ and based on some assumptions it results in finding the function that minimizes the Mean Squared Error (MSE), i.e. finding $\hat{f}(\cdot)$ that minimizes:

$$MSE(f) = \frac{1}{n}\sum_{i=1}^{n}(y_i - \hat{f}(x_i))^2.$$

If we have a homoscedastic error, the average risk of the regression estimator plus a constant value is equal to MSE. *So in small to medium size datasets, leave-one-out or 10-fold cross validation MSE are well-known estimators of the risk function.* For more details on the risk, see Appendix B.

2.3 Local regression methods

Local Polynomial Regression (LPR) assumes that the unknown function $f(\cdot)$ can be locally approximated by a low degree polynomial. Local Polynomial Regression (LPR) fits a low degree polynomial model in the neighborhood (x_i) of the point x. The estimated vector of

parameters used in the fitted LPR is the vector that minimizes a locally weighted sum of squares. Thus for each x a new polynomial is fitted to its neighborhood and the response value is estimated by evaluating the fitted local polynomial with the vector x as covariate. In general the polynomial degree (d) is 1 or 2; for $d = 0$, LPR becomes a kernel regression and when $d = 1$ it changes to LLR.

2.3.1 Definition of Local Polynomial Regression (LPR)

Suppose that the regression function $f(\cdot)$ at the point x can be approximated locally for z inside a neighborhood of x. The idea is to write the Taylor's expansion for z inside a neighborhood of x as follows Fan and Gijbels (1996):

$$f(z) = \sum_{j=0}^{d} \frac{f^j(x)}{j!}(z-x)^j \equiv \sum_{j=0}^{d} \beta_j (z-x)^j. \tag{2}$$

Equation (2) models the regression function by a polynomial function. Thus, for every observation z in the neighborhood of x, we write (2) and estimate the vector $\beta_x = (\beta_{x,0}, \cdots, \beta_{x,d})^T$ by the vector of parameters $\hat{\beta}_x = (\hat{\beta}_{x,0}, \cdots, \hat{\beta}_{x,d})^T$ which minimizes the locally weighted sum of squares defined in Equation (3), and where $\mathcal{K}_b(\cdot)$ represents a kernel function with bandwidth b. In fact, estimating $f(x)$ for the random design as well as for the fixed design results in the locally weighted polynomial regression expressed by the equation below Fan and Gijbels (1996):

$$\hat{\beta}_x = \underset{\beta \in \mathbb{R}^{(d+1)}}{\text{Argmin}} \sum_{i=1}^{n} \mathcal{K}_b(x_i - x) \left(Y_i - \sum_{j=0}^{d} \beta_{x,j}(x_i - x)^j \right)^2 \tag{3}$$

where $\hat{\beta}_x = (\hat{\beta_{x,0}}, \cdots, \hat{\beta_{x,d}})^T$. So the local polynomial estimation for point x is computed as follows:

$$\hat{f}(x) = \sum_{i=1}^{n} a_i(x) Y_i, \tag{4}$$

$$\text{where } a(x) = I_1^T \hat{\beta}_x \text{ and } I_1^T = (1, 0, \cdots, 0).$$

Kernel functions $\mathcal{K}(\cdot)$ are used to weight the observations. They are chosen so that observations closer to the fitting point x have bigger weights and those far from x have smaller weights. If $\mathcal{K}(\cdot)$ is a kernel, then $\mathcal{K}_b(\cdot)$ is also a kernel function.

$$\mathcal{K}_b(u) = \frac{1}{b}\mathcal{K}(\frac{u}{b}), \text{ where } b > 0.$$

Here, b, known as the bandwidth, is a constant scalar value used to select an appropriate scale for the data. In this article, we use the following kernel:

$$\mathcal{K}_b(u) = \frac{1}{b}\mathcal{K}\left(\frac{D(u)}{b}\right), \tag{5}$$

where $D(\cdot)$ is a distance function like the L_2-norm. Some authors including Cleveland (1979) and Cleveland and Devlin (1988), took the K-nearest neighbors of x as its neighborhood. In this case, for each x, $b = D_k(x)$, where $D_k(x)$ is the distance from the K-th nearest neighbors (the farthest neighbor) from the point x. For a detailed discussion see Atkeson et al. (1997).

2.3.2 Bandwidth Selection

A popular bandwidth selection method is the LOO technique suggested in Stone (1977) Stone (1977) which chooses the following bandwidth b:

$$b = \text{Argmin} \sum_{i=1}^{n} (y_i - \hat{f}^{-i}(x_i))^2, \tag{6}$$

where $\hat{f}^{-i}(x_i)$ is the estimation without using the i^{th} observation. Estimating the bandwidth by LOO is a time-consuming task, so it is common to minimize the K-fold cross validation score with $K = 5$ or $K = 10$; this leads to an approximation of LOO. Plug-in bandwidth is another smoothing strategy which is a formula for the asymptotically optimal bandwidth. The plug-in bandwidth requires several unknown quantities that must be estimated from the data. In section 4.2 of Fan and Gijbels (1996), a plug-in bandwidth for linear weighted local linear regression is defined. One of the required parameters for this estimator is $f(\cdot)$'s second derivative which is more difficult to estimate than $f(\cdot)$ itself. In this work we use 10-fold cross validation to find the best bandwidth of our dataset.

2.3.3 Loess

Loess was introduced by Cleveland and Delvin Cleveland and Devlin (1988), and is a multivariate version of LOWESS Cleveland (1979), which is another version of LPR. Loess is described by (4), where the polynomial degree is one $d = 1$ or two $d = 2$. For the bandwidth selection and weight calculation, loess is similar to KNN. Its weights are calculated with (5) where, $u = (x_i - x)$, $D(\cdot)$ is u's L_2-norm in the predictor space and b is the Euclidean distance between the input vector x and its K^{th} nearest neighbor. The weight function chosen by Cleveland and Delvin Cleveland and Devlin (1988) was the Tricube kernel, however it is not mandatory.

In this work, linear loess is the non-parametric smoother function. Therefore, for each input vector x, we use (3), with $d = 1$, to estimate the vector of parameter $\hat{\beta}_x$ from the training set. As we can see in (7), for each prediction the locally weighted linear regression problem is converted to a Weighted Least Squares (WLS) in which the weights are estimated by a kernel function.

$$\hat{\beta}_x = \arg\min \sum_{i=1}^{n} w_i (y_i - x_i^T \beta)^2, \tag{7}$$

where $w_i = \mathcal{K}_b(x_i - x) \geq 0$ is the weight of the i^{th} observation.

2.4 Tolerance intervals

let $\mathbb{X} = (X_1, \cdots, X_n)$ denote a random sample from a continuous probability distribution. A tolerance interval is an interval that is guaranteed, with a specified confidence level γ, to contain a specified proportion β of the population. The $I^T_{\gamma,\beta}$ sign is used to refer to a β-content γ-coverage tolerance interval Krishnamoorthy and Mathew (2009). Then, we have:

$$P_{\mathbb{X}}\bigg(P(X \in I^T_{\gamma,\beta}|\mathbb{X}) \geq \beta\bigg) = \gamma. \tag{8}$$

When our sample set (of size n) comes from a univariate normal distribution, the lower and upper tolerance bounds (X_l and X_u, respectively) are are obtained as follows:

$$X_l = \hat{\theta} - \mathbf{c}\hat{\sigma}, X_u = \hat{\theta} + \mathbf{c}\hat{\sigma} \tag{9}$$

$$\mathbf{c} = \sqrt{\frac{(n-1)(1+\frac{1}{n})Z^2_{1-\frac{1-\beta}{2}}}{\chi^2_{1-\gamma,n-1}}} \tag{10}$$

Where $\hat{\theta}$ is the sample mean of the distribution, $\hat{\sigma}$ its sample standard deviation, $\chi^2_{1-\gamma,n-1}$ represents the $1-\gamma$ quantile of the chi-square distribution with $n-1$ degree of freedom and $Z^2_{1-\frac{1-\beta}{2}}$ is the squared of $(1-\frac{1-\beta}{2})$ quantile of the standard normal distribution Howe (1969).

2.5 Tolerance intervals for least-squares regression

In the case of regression with constant error variance and normal distribution of errors, usually inter-quantiles of a normal distribution with mean zero and variance $\hat{\sigma}^2$, (being the error variance estimator) are used as an approximate solution to find intervals that contain a desired proportion of the distribution of the response variable for a given value of dependent variables. For instance, the 0.95 inter-quantile $[\hat{f}(x) - 1.96\hat{\sigma}, \hat{f}(x) + 1.96\hat{\sigma}]$ is often used as the interval containing 95% of the distribution of $Y(x)$ (i.e., as a regression tolerance interval). As shown by Wallis Wallis (1951), this statement is not true since $\hat{\sigma}^2$ and $\hat{f}(x)$ are only estimations of the true error variance σ^2 and the true mean function at point x, $f(x)$. These estimations are always made on a finite sample and are then pervaded with uncertainty. Tolerance intervals for least squares regression have been introduced in order to take into account this uncertainty. These intervals are described formally by (11). We will refer to such intervals, β-content γ-coverage regression tolerance intervals and they are denoted by $I^T_{\gamma,\beta}(x)$.

$$P\bigg(\int_{U^T_{\beta,\gamma}(x)}^{U^T_{\beta,\gamma}(x)} p_x(t)dt \geq \beta\bigg) = \gamma, \tag{11}$$

where $Y(x) = f(x) + \varepsilon$ and $p_x(t)$ denotes the probability density function of $Y(x)$ for a specified value of the predictor variable x. A two-sided tolerance interval $I^T_{\gamma,\beta}(x)$ for $Y(x)$ is of the form $\hat{f}(x) \pm \rho(x)\hat{\sigma}$, where $\rho(x)$ is the tolerance factor to be determined subject to the content β and the desired confidence level γ. Let $C(x; \hat{f}, \hat{\sigma})$ represent the content of this tolerance interval,

$$C(x;\hat{f},\hat{\sigma}) = P_{Y(x)}\left(\hat{f}(x) - \rho(x)\hat{\sigma} \leq Y(x) \leq \hat{f}(x) + \rho(x)\hat{\sigma}\right). \tag{12}$$

The tolerance factor $\rho(x)$ satisfies the following condition:

$$P_{\hat{f},\hat{\sigma}}\left(C(x;\hat{f},\hat{\sigma}) \geq \beta\right) = \gamma. \tag{13}$$

Equations (11) and (13) could also be expressed as follows:

$$P_{\hat{f},\hat{\sigma}}\left(P_{Y(x)}\left(Y(x) \in I^T_{\gamma,\beta}(x)\right) \geq \beta\right) = \gamma, \tag{14}$$

$$I^T_{\gamma,\beta}(x) = [L^T_{\beta,\gamma}(x), U^T_{\beta,\gamma}(x)] = [\hat{f}(x) - \rho(x)\hat{\sigma}, \hat{f}(x) + \rho(x)\hat{\sigma}].$$

It is important to observe that tolerance intervals in regression are defined separately for each input vector. Therefore, for two different input vectors x_1 and x_2, $I^T_{\gamma,\beta}(x_1)$ and $I^T_{\gamma,\beta}(x_2)$ are different and the event $Y(x_1) \in I^T_{\gamma,\beta}(x_1)$ is independent of $Y(x_2) \in I^T_{\gamma,\beta}(x_2)$. For more details see Hahn and Meeker (1991) and Krishnamoorthy and Mathew (2009).

This intervals has been studied for non-linear parametric in Carroll and Ruppert (1991). They use transformation and/or weighting for the construction of prediction and tolerance intervals for a new response following a non-linear regression fit. Their work addresses the case of normally distributed errors and the non-parametric case in which the error distribution is unknown. *But the non-parametric case concerns the error distribution and not the regression fit, and tolerance intervals for non-parametric regression have not, so far, been applied to non-parametric regression.*

3. State of the Art of the Interval Prediction

Regression models are always built with finite sample size ($n < \infty$), thus the predicted mean or quantile is an estimate of the true unknown conditional mean or quantile of the random variable $Y(x) = f(x) + \varepsilon$. Therefore while dealing with datasets of finite size, we need to make some statistical inferences. This section reviews and gives a comparison of different least-squares and quantile regression techniques used to find such intervals. Besides, we take advantage of this section to address a misunderstood interval prediction method in the machine learning community. We explain its applications and review its drawbacks.

3.1 Interval Prediction Models

The goal of this paragraph is to emphasize the differences between intervals, interval prediction methods and interval prediction models. An interval prediction method is the procedure required for obtaining an interval. It is just the way to do it but when it is applied to a dataset, we obtain an interval prediction model. For example, a tolerance interval for regression is a type of interval. The method to obtain it in linear models is described in Krishnamoorthy and Mathew (2009) and, when applied to a dataset, the model which gives the tolerance interval for each point in the predictor space is the interval prediction model.

Definition 1 *A regression β-content interval prediction model, built on the random sample $\mathcal{S}$, is function $I(\cdot)_\beta$ from the predictor space $\mathbb{R}^p$ to the response variable space $\mathbb{R}$ such that:*

$$I(\cdot)_\beta : \mathbb{R}^p \to \mathcal{I}, \text{ where } \mathcal{I} = \{[a,b]|a,b \in \mathbb{R} \cup \{-\infty,\infty\}, a < b\}. \tag{15}$$

and, the expected content of the intervals is at least β:

$$E_{\mathcal{S}}\left(P\left(Y(x) \in I(x)_\beta \middle| \mathcal{S}\right)\right) \geq \beta. \tag{16}$$

Thus when the *size of our training set goes to infinity* and under certain conditions, a β-content interval prediction model finds intervals that on average contain, at least, a β of the distribution of $Y(x)$. *This is a quite broad definition which covers all the interval prediction method for $Y(x)$ and we will use it for such purpose.*

Our works deals with the regression models, so we omit to mention the regression word and use "interval prediction model" instead of "regression interval prediction model". Note that test and model selection techniques are always applied to models and not to methods. However when a method is more efficient than its competitors on several datasets or in a general theoretical framework, we can state that this method is more efficient than others.

3.2 Least-Squares Based Interval Prediction

We review briefly some well-known statistical inference techniques applied to least-squares regression models. There is an extensive literature on this topic and the goal of this section is to emphasize that least-squares interval prediction methods are not restricted to large sample techniques. However, there are still some subjects like tolerance intervals and simultaneous intervals that need further study. We will see further that prediction and tolerance intervals have some equivalents in the quantile regression set-up, but confidence band and simultaneous tolerance intervals seem to be restricted to the least-squares world.

3.3 Conventional Interval prediction

One of the most common interval prediction techniques used in practice is to take $[\hat{f}(x) - Z_{\frac{1-\beta}{2}}SSE, \hat{f}(x) + Z_{1-\frac{1-\beta}{2}}SSE])$ as the interval which contains a β proportion of $Y(x)$'s population,where SSE^2 is the average MSE given by a LOO or a 10-fold cross validation scheme. One might assume that the intervals expressed below have similar properties to the regression tolerance interval defined in the next section.

$$P\left(Y(x) \in \left[\hat{f}(x) - Z_{\frac{1-\beta}{2}}SSE, \hat{f}(x) + Z_{1-\frac{1-\beta}{2}}SSE\right]\right) = \beta. \tag{17}$$

We assume that:

- the error variance for all x's is constant (homoscedasticity).
- the estimator's variance $\sigma^2_{\hat{f}(x)}$ is constant for all x.
- $\hat{f}(x)$ is an unbiased estimator of $f(x)$.

- the error ε, and $\hat{f}(x)$, are independent and both have normal distributions.

Equation (17) becomes asymptotically valid, but it remains non-applicable for datasets of finite size. For more details on the proof of our statements, refer to Appendix B. *Unfortunately, it is a common practice in the Machine Learning community to employ this method for obtaining intervals having the properties of prediction intervals, tolerance intervals or simultaneous tolerance intervals.* The conventional interval prediction method has some drawbacks:

- First of all, the estimation does not take into account the regression sample size like prediction interval, confidence bands and tolerance intervals.
- It just estimates asymptotic global inter-quantile for the conditional response variable.
- It supposes that the estimated regression function is non-biased.

The problem of finding intervals that satisfy (17) is treated by tolerance for in least-squares regression.

3.3.1 Inference on the conditional mean $f(x)$

This part describes interval prediction techniques that obtain intervals that contain with a confidence level the conditional mean at point x.

- *Point-wise confidence interval for $f(x)$:* The confidence interval for the mean regression function $I_{\beta}^{pw}(x)$ contains asymptotically, a desired proportion β of the conditional distribution of the estimated mean function $\hat{f}(x)$ for each combination of predictors Härdle (1990).
$$P_{\hat{f}}(f(x) \in I_{\beta}^{pw}(x)) = \beta. \tag{18}$$
- *Simultaneous Confidence band for the conditional mean function, for all $x \in \mathcal{X}$:* this is the idea of γ-confidence band $I_{\gamma}^{cb}(x)$ for the regression. These intervals $[U_{\gamma}(x), L_{\gamma}(x)]$ create an envelope around the entire mean regression function $f(\cdot)$ such that, for all $x \in \mathcal{X}$, the probability that the true $f(x)$ is contained in the band is simultaneously γ.
$$P_{\hat{f}}\Big(f(x) \in I_{\gamma}^{cb}(x) \text{ for all } x \in \mathcal{X}\Big) = \gamma, \text{ where } I_{\gamma}^{cb}(x) = [U_{\gamma}(x), L_{\gamma}(x)] \tag{19}$$
where $\mathcal{X}$ is the domain of x.

3.3.2 Inference on the response variable $Y(x) = f(x) + \varepsilon$

This part describes interval prediction techniques that obtain intervals that contain a desired proportion of the conditional distribution of the response variable. These methods are closely related to the methods explained just above.

- *Prediction interval for $Y(x)$:* they are also called expectation tolerance intervals for regression $I_{\beta}^{Pred}(x)$ which contains on average and not at least, a desired proportion

β of the conditional distribution of the response variable $Y(x)$ for each combination of predictors. They are described by the equation below:

$$E_{\hat{f},\hat{\sigma}}\left[P\left(Y(x) \in I_{\beta}^{Pred}(x)|\hat{f},\hat{\sigma}\right)\right] = \beta \text{ where } Y(x) = f(x) + \varepsilon \tag{20}$$

For a detailed discussion about the differences between prediction and tolerance intervals, the reader can find more in Paulson (1943); Hahn and Meeker (1991); Krishnamoorthy and Mathew (2009).

- *Tolerance interval for regression:* tolerance interval for regression $I_{\gamma,\beta}^{T}(x)$ are explained in 2.5. The interval contains, with a confidence level γ, at least a desired proportion β of the distribution of $Y(x)$.

- *Simultaneous tolerance intervals for regression:* β-content γ-coverage simultaneous regression tolerance intervals, described in Krishnamoorthy and Mathew (2009), create an envelope around the entire mean regression function $f(\cdot)$ such that for all $x \in \mathcal{X}$, the probability that $Y(x)$ is contained in the envelope is β, and this coverage is guaranteed with a confidence level γ. They are can be described by (12) and (13) where (13) must be replaced by the equation below:

$$P_{\hat{f},\hat{\sigma}}\left(\min_{x\in\mathcal{X}} C(x;\hat{f},\hat{\sigma}) \geq \beta\right) = \gamma. \tag{21}$$

Note that tolerance intervals and simultaneous tolerance intervals for least squares regression have been well studied for linear regression but the application of these methods in the non-linear and particularly the non-parametric case are limited in the literature.

3.4 Quantile Regression Based Interval prediction

A quantile regression model can estimate one conditional quantile so one-sided and two-sided interval estimation are treated separately.

3.4.1 One-sided interval prediction

- *Quantile regression:* one-sided intervals obtained by the estimation of the conditional quantile are similar to one-sided prediction intervals in least-squares models.

- *Confidence intervals on regression quantiles:* the obtained one-sided interval contains, with a confidence level γ, at least a desired proportion β of $Y(x)$ Kocherginsky et al. (2005); Koenker (1994). They have so far been studied for linear models, and they are similar to one-sided tolerance intervals for regression explained in 2.5.

3.4.2 Two-sided interval prediction

in order to obtain two-sided $(1-\alpha)$-content conditional intervals ($\beta = 1-\alpha$), one must build two distinct quantile regression models: a lower $\frac{\alpha}{2}$-quantile regression model and an upper $(1-\frac{\alpha}{2})$-quantile regression model.

- *Quantile regression:* This is done by a pair of upper and lower quantile regression model. These intervals are estimations and they are similar to two-sided prediction intervals in least-squares models.

- *Confidence intervals on regression quantiles:* These two-sided intervals contain with a γ confidence level, a proportion $1-\alpha$ of the $Y(x)$. As noted, we need a pair of $(\frac{\alpha}{2}, 1-\frac{\alpha}{2})$ quantile regression models but each model now itself needs a one-sided confidence interval on regression quantile. Once we have built the upper and lower quantile regression models, we must obtain a lower (one-sided) $(1-\frac{\tau}{2})$ confidence interval on the lower $\frac{\alpha}{2}$-quantile regression model and an upper (one-sided) $(1-\frac{\tau}{2})$ confidence interval on the upper $(1-\frac{\alpha}{2})$-quantile regression model. By applying the Bonferroni inequality, one can merge the pair of $(1-\frac{\tau}{2})$ confidence intervals to obtain a joint confidence statement with a probability greater or equal to $\gamma = 1-\tau$. More details on this combination can be found in Appendix B.

  It is important to emphasize that although these intervals are theoretically feasible, there has not been any work, until now, which treats the problem of two-sided interval prediction with one-sided confidence intervals on regression quantiles. Such intervals are similar to two-sided γ-coverage $1-\alpha$-content least-squares tolerance intervals and they are explained in 2.5.

As discussed in Koenker (2005), two different quantile regression models may cross or overlap each other, which is called as *quantile crossing*. Thus two-sided interval prediction is more meaningful by enforcing the *non-crossing* constraint. *However after enforcing this constraint, the conditional quantile estimator may not converge to the true conditional quantile. Thus we have to choose between a "non-correct" or non-convergent estimator.*

4. Comparing Interval Prediction Models

For a given dataset, we may have several interval prediction models but we need some quality measure to compare them. For this purpose, we define the dataset measures listed below. These measures are then used as building blocks for some graphical charts and plots explained further in this section. The idea is to provide graphical tools which can help us to compare the effectiveness of different interval prediction methods through different datasets.

4.1 Direct Dataset Measures

For each of the datasets the following quality measures can be computed. Note that the β-content intervals $I(x_i)_\beta = [L(x_i)_\beta, U(x_i)_\beta]$ must be obtained for observations not contained in the training set $\mathcal{S}$. So for small to medium datasets, we obtain these measures with a cross-validation or a LOO schema.

- MIP: Mean Inclusion Percentages: the fraction of the response values that are contained in the β-content interval $I(x_i)_\beta$.

$$MIP_\beta = n^{-1} \sum_{i=1}^{n} V(x_i).$$

where is $V(x_i)$ is defined as:

$$V(x_i) = \begin{cases} 1 & \text{if } Y(x_i) \in I(x_i)_\beta, \\ 0 & \text{otherwise.} \end{cases}$$

- MIS: Mean of Interval Size.

$$MIS_\beta = n^{-1} \sum_{i=1}^{n} size(I(x_i)_\beta).$$

- σ_{is}: sample standard deviation of interval sizes.

$$\sigma_{is} = n^{-1} \sum_{i=1}^{n} (size(I(x_i)_\beta) - MIS_\beta)^2,$$

where $size(I(x)_\beta) = U(x_i)_\beta - U(x_i)_\beta$.

In to verify the reliability of an interval prediction method, we introduce the following constraint.

Definition 2 *Let $\mathcal{S}$ be a sample of size n and $b(n)$ be an function of n, a β-interval prediction model is a reliable model if we have:*

$$\textbf{\textit{MIP Constraint:}}\ MIP_\beta \geq b(n). \tag{22}$$

This definition will further be used in the next section.

4.2 Composed Dataset Measures

We use the above quality measures to define the following composed measures:

4.2.1 Normalized MIS

Suppose that we want to test c different β-content methods ("$Method_1$", "$Method_2$" ,..., "$Method_c$") on the dataset $\mathcal{S}$. They give us c distinct models and each model has a Mean of Interval Size (MIS), so we have: $MIS^1_\mathcal{S}$, $MIS^2_\mathcal{S}$,..., $MIS^c_\mathcal{S}$. But depending on the dataset and β's value, one model may satisfy the MIP constraint or not. For a model that does not satisfy this constraint, we do not compute its normalized MIS. For each model its normalized MIS is equal to the ratio of its MIS to the maximum MIS on this dataset

$$normalizedMIS^i_\mathcal{S} = \frac{MIS^i_\mathcal{S}}{\underset{i\in(1,\cdots,c)}{max}(MIS^i_\mathcal{S})}.$$

If we have :

$$MIS^{m_1}_\mathcal{S} \geq MIS^{m_2}_\mathcal{S} \text{ and } MIP^{m_1}_{\mathcal{S},\beta} \geq b(n) \text{ and } MIP^{m_2}_{\mathcal{S},\beta} \geq b(n)$$
$$\Leftrightarrow m_1 \textbf{ provides a wider reliable envelope than } m_2\textbf{.}$$

M_2 is better than m_1 because it satisfies the MIP constraint and it also gives the smallest normalized MIS value. Choosing the ratio to the maximum MIS value rescales the MIS value between 0 and 1 and lets us compare the strength of methods across different datasets. However we can not use the normalized MIS to compare two models (constructed on the same dataset) that obtain different MIP values but have equal or approximately equal MIS values. In this case, we have to compare them by their Equivalent Gaussian Standard Deviation, explained below.

4.2.2 Equivalent Gaussian Standard Deviation (EGSD)

If we have two reliable models (constructed on the same dataset) having different MIP values but approximately equal MIS values, we normally choose the one that gives the higher MIP. But the situation can get more complicated for models (constructed on the same dataset) with different MIS values and different MIP values. EGSD is a measure which can be used to compare interval prediction models, constructed on the same dataset, which have different MIP values. Such models can have different or equal MIS values. Let m be a β-content interval prediction model built on the dataset $\mathcal{S}$, yielding $MIS^m_{\mathcal{S}}$ and $MIP^m_{\mathcal{S},\beta}$. The idea behind EGSD is to find the Equivalent Gaussian Distribution (EGD) for successful predicted intervals of m. We have seen that by taking intervals size on average equal to $MIS^m_{\mathcal{S}}$, that $MIP^m_{\mathcal{S},\beta}$ of the observations will be contained in their prediction interval. So EGD is the distribution of the size of predicted intervals obtained by model m_1 that correctly contains their response variable. Therefore the EGD which has the smallest variance corresponds to the most efficient model. The Equivalent Gaussian Distribution for m is the normal distribution θ-content inter-quantile size of which will be equal to $MIS^m_{\mathcal{S}}$. We have: $\theta = MIP^m_{\mathcal{S},\beta}$. So the Equivalent Gaussian Standard Deviation of m is calculated by:

$$EGSD^m_{\mathcal{S}} = \frac{MIS^m_{\mathcal{S}}}{2Z_{1-\frac{\alpha}{2}}\theta}, \text{ where } \theta = MIP^m_{\mathcal{S},\beta},\ \alpha = 1-\beta.$$

Now by using each model's EGSD, we can compare models with different values of MIP and MIS. EGSD measures the trade-off between average interval size and the fraction of successful predictions. **Smaller EGSD values denote more effective interval prediction models.** Finally, for the sake of readability, all found EGSD are normalized in each dataset. Thus the final value is the ratio of the method's $EGSD^m_{\mathcal{S}}$ to the maximum $EGSD$ value on the underlying dataset:

$$normalizedEGSD^m_{\mathcal{S}} = \frac{EGSD^m_{\mathcal{S}}}{\underset{i\in(1,\cdots,c)}{max}(EGSD^i_{\mathcal{S}})}.$$

Note that if the model m_1 has smaller EGSD than the model m_2, it does not mean m_2's envelope is wider than m_1's envelope. As seen above smaller normalized MIS values mean smaller envelopes and smaller EGSD values means more effective models.

4.3 Figures

Plots and charts help us to compare different interval prediction methods on different datasets because a figure can visualize complex and big tables. Each plot is dedicated

to one dataset and it compares dataset measures of different interval prediction methods on the same dataset whereas a chart compares a desired dataset measure for different methods and across different datasets. All the presented plots have the same x axis. This axis is labeled "Nominal MIP", and it represents distinct values of the desired proportion (distinct β values). On the other hand, each plot type has a different y axis. This axis denotes the underlying dataset measure on the tested interval prediction models.

4.3.1 MIP plot

The MIP plot is similar to the well-known Q-Q plot with the difference that it compares MIP instead of quantiles. The x axis is denoted by "Nominal MIP" and it represents the desired proportion of inclusion (distinct values of β). The y axis is denoted by "Obtained MIP". It represents the model's MIP. Each point represents the model's obtained MIP for its value on the "Nominal MIP" axis. This figure always has two lines: the "MIP constraint line" and the "Nominal MIP line". The "MIP constraint line" displays $F_{\beta,n}^{0.05}$ for different values of nominal MIP, and the "Nominal MIP line" represents the function $y = x$. By looking at this figure we can see the reliability of a method for different nominal MIP. *The first value in the x axis where a line crosses the MIP constraint line will be called its* ***failure MIP****. It is obvious that the method having the higher failure MIP is the most reliable one.*

One can also use the MIP plot to rate the model's **precision**. If a model obtains MIP values much higher than the desired nominal MIP, it means that the method is reliable but not precise. For example a model which obtains MIP values of 0.45, 0.9 and 0.99 for respective nominal MIP of 0.25, 0.75 and 0.95 is reliable but not precise. The most precise model is the one having the nearest line to the "Nominal MIP line". Finally, the best model in this plot is the one which is the most precise and the most reliable. It means that **the best model in a MIP plot is the one having the nearest line to the upper side of the "Nominal MIP line".**

4.3.2 EGSD plot

The y axis of an EGSD plot is labeled by "Normalized EGSD Value" and it represents the model's normalized EGSD value. By looking at this figure, we can compare the efficiency of different models. It is obvious that the model having the highest line is the most inefficient model. *We suggest using this plot along with the MIP plot to rate the efficiency of reliable methods.* However one may ignore the reliability aspect and take advantage of this plot to compare the efficiency of different models.

4.3.3 MIS plot

The y axis of an EGSD plot labeled "Normalized MIS Value" and it represents the model's normalized MIS value. By looking at this figure, we can compare the model which obtains the tightest reliable envelope. The model having the highest line provides the widest envelope. If a model does not satisfy the MIP constraint, its normalized MIS value is not computed. The MIS plot shows each model's normalized MIS until its "failure MIP". *We suggest using this plot along with the EGSD plot.*

4.3.4 Charts

Charts are used to compare one dataset measure on different datasets. We propose the following charts:

- Mean Inclusion Percentage Chart (MIP chart): the goal of this chart is to compare the mentioned methods based on their fraction of response values located inside their predicted intervals. It just displays the MIP value and it usually does not contain important information.
- MIS ratio chart: this chart displays the normalized MIS measure on different datasets.
- Equivalent Gaussian Standard Deviation chart (EGSD chart): it displays the normalized EGSD measure on different datasets.

5. Predictive Interval Framework

The goal of this section is to propose a new interval prediction framework. We introduce the concept of regression predictive intervals and regression Predictive Interval Model (PIM). Next, we propose a statistical test to verify if an "interval prediction model" is a "predictive interval model". In the same context, we introduce two measures for rating interval prediction models. These measures rate the efficiency and the tightness of the obtained envelope.

5.1 Predictive Interval Models (PIMs)

In the frequentist interpretation of confidence intervals, a confidence interval for a parameter contains zero or one parameter. The parameter is fixed and confidence intervals change with different random samples. In the same way, the confidence level (γ) used in tolerance intervals for regression defined in (14) and confidence intervals for regression quantiles (explained in Appendix B) mean the following: the probability that the obtained intervals contain, under re-sampling, at least a proportion β of the conditional distribution of the response value $Y(x)$ is γ. We know that the confidence level in Neyman-Pearson confidence intervals is independent of the observed sample. It means that if we obtain γ-confidence β-content tolerance intervals of an observed sample from a regression function, then the confidence level γ does not induce any posterior probability of including β proportion of the distribution of $Y(x)$. Therefore, the confidence coefficient in frequentist confidence intervals cannot be interpreted as posterior probability. This idea is discussed in detail in Chapter 7 in Walley (1991).

Hence, under the frequentist viewpoint of regression, the true conditional response variable's inter-quantile is included with probability zero or one in the obtained interval (by using tolerance intervals for regression or confidence intervals for regression quantiles). Our goal is to obtain intervals that correctly bracket these inter-quantiles. They can be found in two ways: the first approach takes a very high confidence level like $\gamma \approx 1$ and the second method finds the smallest confidence level $0 < \gamma_0 < 1$ which includes the true unknown model. We introduce the concept of predictive intervals which refer to both of these intervals. A predictive interval built on $\mathcal{S}$, is guaranteed to contain for the query point x,

at least a desired proportion of the conditional distribution of the response variable. *It can be obtained with tolerance intervals for regression or confidence intervals for regression quantiles but these concepts have so far only been treated for linear models*

Definition 3 *Let $\mathcal{S} = \{(x_1, Y_1) \cdots, (x_n, Y_n)\}$ denote a random sample where $Y_i = f(x_i) + \varepsilon_i$ and ε_i is white noise. A β-content predictive interval for x, denoted by $I(x)^P_\beta$, is an interval such that:*

$$P_{Y(x)}\left(Y(x) \in I(x)^P_\beta \middle| \mathcal{S}\right) \geq \beta, \text{ where } I(x)^P_\beta = [L(x)^P_\beta, U(x)^P_\beta]. \tag{23}$$

Since we have observed $\mathcal{S}$, $I(x)^P_\beta$ is no longer random and the probability measure is just related to cover at least a proportion β of the conditional distribution of the response variable $Y(x)$ for a specified combination of the predictors.

Definition 4 *Let $\mathcal{S} = \{(x_1, Y_1) \cdots, (x_n, Y_n)\}$ denote a random sample where $Y_i = f(x_i) + \varepsilon_i$ and ε_i is white noise. A β-content predictive interval model, denoted by $I(\cdot)^P_\beta$, is a function such that:*

$$I(\cdot)^P_\beta : \mathbb{R}^p \to \mathcal{I}, \text{ where } \mathcal{I} = \{[a, b] | a, b \in \mathbb{R} \cup \{-\infty, \infty\}, a < b\}, \tag{24}$$

and for all x in the predictor space, the obtained interval is a β-content predictive interval described by (23).

5.2 Predictive Interval Model Test (PIM Test)

The goal of this part is to develop a statistical test with which we can rate the reliability of any interval prediction model claiming to provide β-content predictive intervals. A predictive interval model must provide predictive intervals for each point in the predictor space. We saw that the distribution of $Y(x)$ changes for each value of x. So, in order to see whether an interval for the regression function at the point x contains at least a proportion β of the s distribution of $Y(x)$, we need (for each combination of predictors x) a sample set from the distribution of $Y(x) = f(x) + \epsilon$, and then we can observe if the constructed interval contains a proportion s distribution of the distribution of $Y(x)$. In the same way, in order to verify if the methods works for an entire dataset $\{x_i | \in (1, \cdots, n)\}$, we need a distinct sample set for each x_i and this sample must be drawn from $Y(x_i)$. Since a sample set is required for each $x_i, i \in (1, \cdots, n)$, the described procedure requires a huge dataset having many observations for each point x in the feature space which makes it impractical or impossible for multiple regression problems. However, we can make some approximations and use the results stated above to derive the following test. We first begin by defining a variable obtaining MIP on the dataset. Then we will see that this variable can be approximated by a normal distribution, so we use the quantiles of the standard normal distribution as the function $b(n)$ used in Equation (22).

5.2.1 Simultaneous Inclusion for Predictive Intervals

A β-content predictive interval $I(x)^P_\beta$ must contain at least β proportion of the conditional distribution of $Y(x)$. Hence, the probability in (23) is just related to contain at least a proportion β of the conditional distribution of the $Y(x)$. We define the function $V(x)$ as:

$$V(x) = \begin{cases} 1 & \text{if } Y(x) \in I(x)^P_\beta, \\ 0 & \text{otherwise.} \end{cases}$$

The above definition means that the probability that $V(x)$ is equal to 1 is β, so $V(x)$ has a Bernoulli distribution with $p = \beta$.

$$V(x) \sim Bernoulli(\beta). \tag{25}$$

Suppose that we have a dataset of n_{train} observations $\mathcal{T} = \{(x_1, Y_1) \cdots, (x_{n_{train}}, Y_{n_{train}})\}$ with which we build our model, and n_v other observations $\mathcal{S} = \{(x^v_1, Y^v_1) \cdots, (x^v_{n_v}, Y^v_{n_v})\}$, not contained in the original dataset as our test set. If we apply the function $V(\cdot)$ on the whole test set $\mathcal{S}$ and sum the result, we obtain:

$$MIP_\beta = n_v^{-1} \sum_{i=1}^{n_v} V(x^v_i). \tag{26}$$

Therefore, we can deduce that $MIP_{\mathcal{S},\beta}$ has a Binomial distribution. This is expressed formally by (27) where $Binomial(n_v, \beta)$ is a binomial distribution with $n = n_v$ and $p = \beta$.

$$n_v MIP_\beta \sim Binomial(n_v, \beta). \tag{27}$$

If n_v is sufficiently large, we can assume that MIP_β has a normal distribution as:

$$MIP_\beta \sim \mathcal{N}(\beta, \frac{\beta(1-\beta)}{n_v}). \tag{28}$$

Thus in a PIM, the fraction of instances having their response value included in their predictive intervals is on average β. This means such predictive intervals for regression have in average a simultaneous content of β so, on average, they are like simultaneous regression tolerance intervals. **For small to medium datasets, MIP_β is computed in a cross-validation or leave-one-out schema on the whole dataset which means that $\mathcal{S} = \mathcal{T}$.**

5.2.2 Testing Predictive Interval Models

As we have seen in (28), the random variable MIP_β can usually be well approximated by a normal distribution. The test below is used to verify, with level α, if the interval prediction method does provide β_0-content predictive intervals for all x in the predictor space.

$$H_0 : \beta \geq \beta_0 \ \textit{versus}\ H_1 : \beta < \beta_0, \tag{29}$$

then H_0 can be rejected with significance α where:

$$MIP_\beta < n_v^{-1/2} \beta_0(1-\beta_0) Z_\alpha = F^\alpha_{\beta_0, n_v}, \tag{30}$$

where Z_α is the α-quantile of the standard normal distribution. So if (30) is not true, we fail to reject the null hypothesis with significance level α, and accept the underlying model

as a model providing β-content predictive intervals.

In this work we used a significance level of $\alpha = 0.05$, so for each dataset we compared the MIP_β's value on the test set with $F^{0.05}_{\beta_0,n_v}$. Thus, any PIM must pass the PIM test as defined below:

$$\textbf{PIM Test: } MIP_\beta \geq F^{0.05}_{\beta_0,n_v}. \tag{31}$$

For the sake of simplicity, we refer to $MIP_{\mathcal{S},\beta}$ for a given dataset $\mathcal{S}$ and desired proportion β as MIP (Mean Inclusion Percentage). As we have seen in (28), the fraction of response values inside their β-content predictive intervals converges to β, so the test defined in (31), where $\alpha = 0.05$ is used to verify if the obtained intervals, with a confidence level 0.95 *and on average and not at least like in simultaneous tolerance internals*, do simultaneously contain a proportion β_0 of the distribution of $Y(x)$ for all x in the predictor space.

Remark

The average simultaneous content of PIMs mentioned just above, means that: for a PIM built on the training set $\mathcal{T}$, if we test the PIM with a large number of new observations $(x_i^*, Y(x_i^*))_{i\in\{1,\cdots,n_v\}}$, the fraction of the new response values included in their predictive intervals **is guaranteed to be** β. This is expressed formally in (28).

Note that, this is not the same as the average content for any β-content interval prediction models. A β-content interval prediction model built on a specified training set $\mathcal{T}$, **is not guaranteed to contain**, on average, a β proportion of successful predictions. However, if we have a large number l of β-content interval prediction model $I(\cdot)_{\mathcal{T}_j,\beta}$, built on a large number l of training set $(\mathcal{T}_j)_{j\in\{1,\cdots,l\}}$, and we test each $I(\cdot)_{\mathcal{T}_j,\beta}$ with a large number of unseen instances $(x_i^*, Y(x_i^*))_{i\in\{1,\cdots,n_v\}}$. In this case, the expected probability of each β-content interval $I(x_i^*)_{\mathcal{T}_j,\beta}$ to contain a proportion β of the distribution of $Y(x_i^*)$ is β. This is expressed formally in (16).
Thus, having a β-content interval prediction model built on the training set $\mathcal{T}$, if we test the model with **large number of training set and new testing observations,** the fraction of response value included in their predictive intervals is on average β. An arbitrary β-content interval prediction method needs a large number of training set and new testing observations, or one large training set and a large number of new testing observations, to guarantee an MIP_β equal to β. However a PIM, must guarantee an $MIP_\beta = \beta$ with any training set.

5.3 Hyper-parameter Tuning and Model Selection

This part addresses hyper-parameter tuning questions related to predictive interval explained in Section 5.1. We first convert the hyper-parameter tuning problem to an optimization problem. Then, we propose a simple optimization algorithm that takes advantage of existing least-squares regression method to find an efficient model for predicting both the mean and the predictive interval. Finally, we give an application with tolerance intervals for least-square regression. Note that, the dataset $\mathcal{S}$ does not change in the hyper-parameter tuning phase, so we do not use it to index our variables.

5.3.1 The Optimization Problem

Let λ denote the vector of hyper-parameter of the β-content interval prediction method $I(\cdot)_{\mathcal{S},\beta}$. The vector λ is divided into two set of elements. The first set is the vector of hyper-parameter for the regression method and the second set is specific to the interval prediction method. Our goal is to find the optimal value of λ, for any β-content interval prediction model that is also a β-content predictive interval model (pass the PIM test). This result in the following problem:

$$\lambda_0 = \text{Argmin}(MIS_\beta^\lambda), \text{ where } MIS_\beta^\lambda = \frac{1}{n}\sum_{i=1}^{n} I(x_i)_\beta^\lambda \tag{32}$$

Subject to:

$$\begin{aligned} &\textbf{PIM Tuning Constraint: } MIP_\beta^{\lambda_0} = \beta \\ &\lambda\textbf{-Specific Constraints: } \text{depends on the PIM.} \end{aligned} \tag{33}$$

Note that the MIP constraint is a hard constraint and there is no trade-off between satisfying this constraint and minimizing the MIS. We found λ_0 which satisfies the constraint defined above. This results in intervals having the smallest MIS measure where MIP and MIS are computed based on a leave-one-out or 10-fold cross validation scheme on the training set.

5.3.2 The Optimization Algorithm

We propose the following hyper-parameter tuning method:

- First, find the regression model with the smallest prediction error MSE.
- Then use an optimization method that finds the λ_0 that satisfies the tuning constraint and also results in intervals having the smallest MIS measure on the previously found regression model.
- Note that the space of λ in the second step is restricted to λ values that have their regression hyper-parameter equal to the hyper-parameter of the regression model found in the first step.

This process divide the optimization in two phase, the hyper-parameter of the regression method tuning and the hyper-parameter tuning of the interval prediction method. In order to obtain the smallest intervals (measured by MIS), we need the smallest values of prediction errors. So, choosing the regression model that minimizes the LOO or 10-fold cross validation MSE can give the most efficient PIM. This approach does not always find the optimal solution but remains a simple and efficient tuning algorithm. By using this algorithm, we can take advantage of existing regression packages to find the best regression model, and then use the aforementioned method to tune the the remaining hyper-parameters. Finally, this approach give us one model for predicting both the mean and the predictive interval.

5.3.3 An Application with Tolerance Intervals

We know that β-content predictive intervals can be obtained via tolerance intervals for regression or via confidence interval on regression quantiles and such intervals are obtained upon regression models which may themselves have hyper-parameters (For the sake of brevity we continue this part with tolerance intervals for regression but the same procedure and statements hold for confidence interval on regression quantiles). Consider an example of constructing Predictive Interval Model (PIM)s by tolerance intervals on a KNN regression. This model has two hyper-parameters, the KNN regression's hyper-parameter which is the number K, and the confidence level γ which is the coverage level in γ-coverage β-content tolerance intervals. So we have $\lambda = (K, \gamma)^T$. First we find the regression's hyper-parameters K; it is K for KNN but it can be kernel related parameters in SVM regression or nothing in the linear regression (this hyper-parameter depends on the regression method). Once we have found the best regression model, we use an iterative algorithm that searches the smallest γ that satisfies the tuning constraint defined above. High values of γ will guarantee the PIM tuning constraint but the computed intervals can be very large, so the search begins with a high confidence value like $\gamma = 0.9$ or $\gamma = 0.99$ and we try to decrease γ and thus decrease the mean interval size. This procedure is repeated as long as the tuning constraints are satisfied and the search strategy is left to the user. Some datasets might require just 2 or 3 iterations but some others may work with small γ. It depends on the dataset and it can be influenced by the domain expert.

6. Tolerance Intervals for Local Linear Regression

In the previous section we introduced the predictive interval framework. In this section we introduce two methods for obtaining tolerance intervals for local linear regression which in the next section, will be employed to obtain PIMs for LLR. We assume that the mean regression function is locally linear and the prediction error is locally homoscedastic. Our methods do not neglect the regression bias and find variable size intervals that work properly with biased regression models. The proposed tolerance intervals are constructed based on the leave-one-out or 10-fold cross validation prediction errors of the local linear regression. The local linear regression needs a regression bandwidth which could be found by any of the existing methods in the literature. In order to obtain our non-parametric predictive intervals, we need a second bandwidth, which is the tolerance interval bandwidth (LHNPE bandwidth). This work suggests two different tolerance interval bandwidths: a bandwidth having a fixed number of neighbors and a bandwidth having a variable one but both obtain variable size intervals.

Another important subject is the regression bias. It is well known that the optimal smoothing non-parametric regression method consists of balancing between the bias and the standard deviation. This non-parametric bias does not vanish even with large sample sizes, so it is important for interval prediction methods to take it into account. The idea behind our tolerance intervals is to exploit the local density of prediction error $(Y_i - \hat{f}(x_i))$ in the LHNPE neighborhood (explained further) of the x^* to find the most appropriate intervals that contain the desired proportion of the distribution of $Y(x^*)$. We find tolerance intervals for the response variable by adding the regression estimates to the tolerance intervals for the prediction error. Prediction error's tolerance intervals are centered on the estimation of

negative bias, so when added to the regression estimates, they remove the regression bias. Therefore, we obtain response variable's tolerance intervals which correctly contains, with confidence γ, a proportion β of $Y(x^*)$.

6.1 Theoretical Aspect

This part describes the theoretical context of tolerance intervals for local linear regression. We first define the concept of a Local Homoscedastic Normal Prediction Error (LHNPE) regression estimator. Then, we define the LHNPE neighborhood of a query point and in the end we will use a simple straightforward inference to obtain the formula of tolerance intervals for local linear regression.

Definition 5 *The oscillation of the function $f : X \to \mathbb{R}$ on an open set U is defined as:*

$$\omega_f(U) = \sup_{x \in U} f(x) - \inf_{x \in U} f(x).$$

Definition 6 *A regression estimator $\hat{f}(x)$ is a **Local Homoscedastic Normal Prediction Error (LHNPE)** if it satisfies the following conditions:*

- *Normal distribution: the prediction error $\varepsilon_x^{pred} = Y(x) - \hat{f}(x)$ has a normal distribution.*
- *Almost constant distribution the prediction error: We suppose that the mean $\mu(\varepsilon_x^{pred})$ and the standard deviation $\sigma(\varepsilon_x^{pred})$ of the distribution for the prediction error have small local oscillations. This is defined formally as:*

  For all x, there exists an open set $U \ni x$, such that:

  $$\omega_{\mu(\varepsilon_x^{pred})}(U) \leq \upsilon_1 \text{ and } \omega_{\sigma(\varepsilon_x^{pred})}(U) \leq \upsilon_2,$$

  where υ_1 and υ_2 are small fixed positive values.

Definition 7 *Let $\hat{f}(x^*)$ be a LHNPE regression estimator for the query point x^*. The **LHNPE neighborhood for** x^* are instances for which the prediction error satisfies the LHNPE conditions. This neighborhood is described as below:*

$$Kset_{x^*} = \{(x_i, Y_i) | d(x^*, x_i) \leq b\}, \tag{34}$$

where $d(x^*, x_i)$ is a distance function in the feature space and b denotes the LHNPE bandwidth.

Note that the LHNPE bandwidth $Kset_{x^*}$ is different from the regression bandwidth Reg_{x^*} in local linear regression:

$$Reg_{x^*} = \{(x_i, Y_i) | d(x^*, x_i) \leq b_{reg}\}. \tag{35}$$

The regression bandwidth minimizes the regression bias-variance trade-off but the LHNPE bandwidth is used to find the neighborhood which satisfies the LHNPE conditions. The LHNPE neighborhood is almost always included in the regression neighborhood:

$$Kset_{x^*} \subseteq Reg_{x^*}, \tag{36}$$

because the constant's $\left(Y(x^*) - \hat{f}(x^*)\right)$ distribution in the neighborhood of the query point x^* usually occurs inside its regression neighborhood. It is possible to find two different regression neighborhoods being next to each other having approximately the same prediction error distribution and not the same regression neighborhood. There are already several references on regression bandwidth Reg_{x^*} selection in non-parametric regression. We do not treat this problem and the reader can find more details in Fan and Gijbels (1996) and Härdle (1990).

Proposition 1 *Let $Y(x) = f(x) + \varepsilon_x$ denote a regression function and let $\hat{f}(x)$ denote its Local Linear regression estimator. If our regression estimator satisfies the conditions below:*

- *Normal error distribution: $\varepsilon_x \sim \mathcal{N}(0, \sigma_x^2)$.*
- *Normal distribution of the local linear estimator: $\hat{f}(x) \sim \mathcal{N}\left(f(x) + Bias_{\hat{f}(x)}, \sigma^2_{\hat{f}(x)}\right)$, Fan et al. (1995) have shown that this assumption holds under certain regularity conditions.*
- *$\hat{f}(x)$ satisfies the LHNPE conditions defined above.*

where $Bias_{\hat{f}(x^*)} = E[\hat{f}(x^*) - f(x^*)]$ is the estimator's bias, σ_x^2 is the variance of the error and $\sigma^2_{\hat{f}(x^*)}$ is the variance of the estimator. Then the γ-confidence β-content regression tolerance interval for the query point x^* is:

$$I(x^*)^T_{\gamma,\beta} = \hat{f}(x^*) + I(\varepsilon^{pred}_{x^*})^T_{\gamma,\beta}, \tag{37}$$

$$\text{where } \varepsilon^{pred}_{x^*} = Y(x^*) - \hat{f}(x^*).$$

In the above equation, $I(x^*)^T_{\gamma,\beta}$ and $I(\varepsilon^{pred}_{x^*})^T_{\gamma,\beta}$ denote, respectively, the regression tolerance interval and the prediction error tolerance interval.

Proof: See Appendix B.

Even though we have a biased prediction, our tolerance interval for $Y(x^*)$ contains the desired proportion of the conditional distribution of the response variable. This is due to the fact that our tolerance intervals on the response variable $I(x^*)^T_{\gamma,\beta}$ are computed based on the tolerance intervals on the prediction error $I(\varepsilon^{pred}_{x^*})^T_{\gamma,\beta}$. LHNPE conditions assume that the prediction error has an unknown normal distribution with mean and variance being respectively the negative bias and the variance of the prediction error. So, for high values of γ and for $\beta > 0.5$, $I(\varepsilon^{pred}_{x^*})^T_{\gamma,\beta}$ will contain the true bias. Therefore, adding $I(\varepsilon^{pred}_{x^*})^T_{\gamma,\beta}$ to the biased regression estimate will remove the bias and give tolerance intervals that works properly with biased regression estimators.

6.2 Computational Aspect

By taking advantage of the LHNPE conditions for the local linear estimator, the tolerance interval on the prediction error at the point x^*, described by (37), is approximated by the tolerance interval on prediction errors inside its LHNPE neighborhood. The prediction

error inside the LHNPE neighborhood of the query point is represented by $Eset_{x^*}$ and it is defined formally as:

$$Eset_{x^*} = \{\varepsilon_i^{pred} | (x_i, Y_i) \in Kset_{x^*}\}, \text{ where } \varepsilon_i^{pred} = Y_i - \hat{f}^{-i}(x_i), \tag{38}$$

where $\hat{f}^{-i}(x_i)$ is the local linear estimation without using the i^{th} observation, obtained by (4). Note that when (x_i, Y_i) is in our training set, $Y_i - \hat{f}(x_i)$ becomes a residual and it depends on the random variable Y_i; however, $Y_i - \hat{f}^{-i}(x_i)$ and Y_i are independent.

Hence, given an input vector x^*, K the number of neighbors in $Eset_{x^*}$, β the desired content and γ the confidence level, the tolerance interval for the prediction error variable $\varepsilon_{x^*}^{pred}$ is computed by replacing $\hat{\theta}, \hat{\sigma}$ and n in Equations (9) and (10) which results in:

$$I(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T} = \hat{\theta} \pm \mathbf{c}\hat{\sigma}, \text{ where } \mathbf{c} = \sqrt{\frac{(K-1)(1+\frac{1}{K})Z^2_{1-\frac{1-\beta}{2}}}{\chi^2_{1-\gamma,K-1}}}, \tag{39}$$

$$\hat{\theta} = \bar{\varepsilon}_i^{pred} = K^{-1} \sum_{\varepsilon_i^{-i} \in Eset_{x^*}} \varepsilon_i^{pred} \text{ and } \hat{\sigma}^2 = (K-1)^{-1} \sum_{\varepsilon_i^{pred} \in Eset_{x^*}} (\varepsilon_i^{pred} - \bar{\varepsilon}_i^{pred})^2. \tag{40}$$

We propose to take the LHNPE neighborhood as the K-nearest neighbors to the query points where K can be a fixed or a variable number tuned on the dataset. So depending on the LHNPE neighborhood selection method, we have two different methods to obtain tolerance intervals for LLR but both methods require 10-fold cross validation or LOO errors of the whole training set. We denote this by *error_set*:

$$error\_set = \{\varepsilon_i^{pred} | (x_i, Y_i), i \in (1, \cdots, n)\}, \text{ where } \varepsilon_i^{pred} = Y_i - \hat{f}^{-i}(x_i). \tag{41}$$

Algorithm 1 summarizes the required steps for obtaining tolerance intervals for local linear regression.

Algorithm 1 Tolerance Interval for local linear regression

1: **for all** $(x_i, Y_i) \in trainingSet$ **do**
2: $\quad \varepsilon_i^{pred} \leftarrow Y_i - \hat{f}^{-i}(x_i)$
3: $\quad error\_set \leftarrow \{error\_set, \varepsilon_i^{pred}\}$
4: **end for**
5: **for all** $x^* \in testSet$ **do**
6: $\quad fval \leftarrow \hat{f}(x^*)$
7: $\quad Kset_{x^*} \leftarrow$ findToleranceNeighborhood(x^*)
8: $\quad Eset_{x^*} \leftarrow$ error of instances in $Kset_{x^*}$, previously stored in *error_set*
9: $\quad I(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T} \leftarrow \beta$-content γ-coverage normal tolerance interval of $Eset_{x^*}$ as in Equations (39,40).
10: $\quad I(x^*)_{\gamma,\beta}^{T} \leftarrow fval + I(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T}$
11: **end for**

6.3 LHNPE bandwidth with Fixed K

In this method we take a fixed K number of the nearest neighbors of x^* as its LHNPE neighborhood. These neighbors are returned by the function "findToleranceNeighborhood(x^*)".

K is a hyper-parameter and is tuned on the training set. We denote this interval prediction method for LLR by "Fixed K". Once the local linear model has been built and *error_set* has been found on the training set, the computational complexity of interval prediction for a new instance is the same as an evaluation under the local linear regression. We select this neighborhood in such a way that it remains inside the regression neighborhood. This condition is respected appropriately by all points of the feature space of a dataset. Thus we have to take a LHNPE bandwidth that is coherent on the majority of points in the feature space. In our experiments, this condition is always satisfied except in the "Auto" dataset where the LHNPE bandwidth is a bit greater than the Regression bandwidth.

6.4 LHNPE bandwidth with Variable K

The idea behind this LHNPE bandwidth selection method is to find the "best" LHNPE bandwidth (best K) of each input vector x^*. This method is summarized in Algorithm 2. For a fixed value of β, and for each input vector x^*, the computation begins with an initial value of K, then the β-content γ-coverage normal tolerance interval of errors in $Eset_{x^*}$ defined in (38) is calculated. This process is repeated for the same input vector x^* but different values of $K, MIN_K \leq K \leq MAX_K$. Finally, the $I(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T}$ having the smallest size among the tolerance intervals computed by different values of K (different $Eset_{x^*}$) is chosen as the desired interval and is added to $\hat{f}(x^*)$. This iterative procedure leads us to choose the interval that has the best trade-off between the precision and the uncertainty to contain the response value. The more K increases, the less the local homoscedasticity assumptions match reality and this yields a prediction error variance different from the true one. If we find a variance higher than the true one, it could be partially compensated by the fact that the tolerance interval size decreases when the sample size increases. However, an increase in K may lead us to obtain smaller prediction variance; this issue is controlled by MAX_K. On the contrary, when K is small, the LHNPE conditions are respected but the tolerance interval sizes increase just because the sample size is too small. Thus choosing the value of K that minimizes a fixed β-content γ-coverage tolerance interval ensures that we will have the best trade-off between the faithfulness of the local assumptions (LHNPE conditions) and the required sample size to guarantee the desired β proportion of the response value. The optimal value of K may vary much more on heterogeneous datasets.

MIN_K and MAX_K are global limits for the search process. MAX_K stops the search process if the best value for K is not found before. This can occur when increasing the neighborhood, it gets contaminated with instances having smaller prediction errors than the prediction of the query point. In practice, these smaller prediction errors usually belong to a different subspace of the feature space with different error variances and/or prediction error distributions. Therefore these two bounds serve to restrict the search process in a region where it is most likely to contain the best neighborhood of x^*. MAX_K is usually included in the regression neighborhood. However one can take it greater than the regression bandwidth and let our search algorithm (Algorithm 2) find the neighborhood which gives the smallest tolerance interval.

Once the local Linear model has been built and *error_set* has been found on the training set, the computational complexity of interval prediction for a new instance is ($MAX_K -$

Algorithm 2 LHNPE neighborhood with variable K

function findToleranceNeighborhood(x^*)
  $IntervalSize_{min} \leftarrow \infty$
  $Kset_{return} \leftarrow \emptyset$
  for all $i \in MIN_K, \ldots, MAX_K$ **do**
    $Kset_{x^*} \leftarrow i$ nearest number of instances $(x_i, Y_i) \in trainingSet$ to x^*
    $Eset_{x^*} \leftarrow \varepsilon_i^{-i}$ of instances in $Kset_{x^*}$ previously computed in $error\_set$
    $I(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T} \leftarrow \beta$-content γ-coverage normal tolerance interval of $Eset_{x^*}$ as in Equations (39,40).
    if $size(I(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T}) \leq IntervalSize_{min}$ **then**
      $Kset_{return} \leftarrow Kset_{x^*}$
      $IntervalSize_{min} \leftarrow size(I(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T})$
    end if
  end for
  return $Kset_{return}$
end function

MIN_K) times higher than the complexity of an evaluation under the local linear regression. Because from the beginning to the $Kset_{x^*}$-finding step, everything is similar to LLR, then in the interval calculation phase, LLR computes just one value and "Var K." computes $(MAX_K - MIN_K)$ intervals. More explanation on the LLR complexity can be found in Section 2.3

7. Predictive Interval Models for Local Linear Regression

This section describes how to use tolerance intervals for local linear regression to obtain predictive interval models. First, we describe how the confidence level γ in these tolerance intervals can be used to obtain predictive interval models. Then, we see how to find the "best" value of γ that provides predictive interval model with the smallest mean interval size.

7.1 The General Formula

The β-content predictive interval on the prediction error, denoted by $I(\varepsilon_{x^*}^{pred})_{\beta}^{P}$, is obtained by finding the predictive intervals hyper-parameters which satisfies the PIM tuning constraint in (33). Finally, the β-content predictive interval on the response variable is computed by adding local linear regression estimation to the error predictive interval:

$$I(x^*)_{\beta}^{P} = \hat{f}(x^*) + I(\varepsilon_{x^*}^{pred})_{\beta}^{P}.$$

As explained in 5.3, regression predictive intervals models have two types of hyper-parameters. This first is the regression method's hyper-parameter. In LLR, it is the bandwidth used for regression and it serves to find the $error\_set$. The second type of hyper-parameters are the predictive interval hyper-parameters. These hyper-parameters are (K, γ) or (MIN_K, MAX_K, γ), respectively, for predictive intervals with fixed K and predictive intervals with variable K.

7.2 Application with Linear Loess

We saw above, how to compute predictive interval models in the general form of local linear regression. This paragraph briefly reviews an application with the linear loess regression method. Loess is a version of linear polynomial regression that for each query point, takes its K nearest instances in the feature space as its neighborhood. We denote loess's regression bandwidth with K_{loess}. Loess uses a first or second degree polynomial, so Linear loess refers to a loess with a first degree polynomial.

Predictive intervals with Linear loess have three or four hyper-parameters: the linear loess bandwidth K_{loess} and the predictive hyper-parameters which are the confidence level γ and the LHNPE bandwidth. As seen above, (K) and (MIN_K, MAX_K) are respectively the LHNE bandwidth for predictive intervals with fixed K and predictive intervals with variable K. Based on (36), we usually have :

$$MAX_K \leq K_{loess} \textbf{ or } K \leq K_{loess}.$$

7.3 Hyper-parameter Tuning

At this stage, we suppose that the loess bandwidth K_{loess} has been found. The only difference is that in variable K, we are looking for the pair (MIN_K, MAX_K) instead of K in fixed K. So we have $\lambda = (K_{loess}, (\gamma, MIN_K, MAX_K))$ for variable K and $\lambda = (K_{loess}, (\gamma, K))$ in fixed K. The tuning procedure explained here is similar to the one discussed in 5.3. *So once we have obtained the regression model, the PIM hyper-parameter tuning reduces to the constraint optimization problem listed below where* ***all the constraints are hard constraints****.*

Optimization problem for fixed K:

$$(\gamma, K) = \text{Argmin}(MIS_\beta^\lambda), \text{ where } MIS_\beta^\lambda = \frac{1}{n}\sum_{i=1}^{n} I(\varepsilon_{x_i})_{\gamma,\beta}^T$$

$$\text{Subject to } \textit{Tuning Constraints: } \begin{cases} \textit{PIM Tuning Constraint: } MIP_\beta^{\lambda_0} = \beta \\ \lambda\textit{-specific Constraints:} \begin{cases} 0 < \gamma < 1 \\ 0 < K \leq n \end{cases} \end{cases}$$

Optimization problem for variable K:

$$(\gamma, MIN_K, MAX_K) = \text{Argmin}(MIS_\beta^\lambda), \text{ where } MIS_\beta^\lambda = \frac{1}{n}\sum_{i=1}^{n} I(\varepsilon_{x_i})_{\gamma,\beta}^T$$

$$\text{Subject to } \textit{Tuning Constraints: } \begin{cases} \textit{PIM Tuning Constraint: } MIP_\beta^{\lambda_0} = \beta \\ \lambda\textit{-specific Constraints:} \begin{cases} 0 < \gamma < 1 \\ 0 < MIN_K \leq MAX_K \leq n \end{cases} \end{cases}$$

Algorithm 3 describes how to tune the predictive interval hyper-parameters for variable K. The algorithm used for the fixed K is almost the same so we do not mention it. In

a first attempt, γ is considered as a fixed high value like $\gamma = 0.9$ or $\gamma = 0.99$ and we focus on finding the LHNPE neighborhood hyper-parameter: the hyper-parameter K or the pair (MIN_K, MAX_K). We saw that the variable MIP_β^λ defined by Equation (33) must on average be greater than or equal to $n\beta$. Thus we can select the LHNPE neighborhood hyper-parameter(s) which find(s) intervals that, based on a LOO or 10-fold cross validation scheme on the training set, satisfies the tuning constraint defined in (33) and also have the smallest Mean Interval Size (MIS). Once we have found K or (MIN_K, MAX_K) we search for the smallest value of γ that satisfies the PIM tuning constraint.

Small neighborhoods result in big tolerance intervals, thus higher coverage. As long as K's value is increased, the mean interval size decreases too. However after a threshold, the mean interval size may increase or change a little bit but the coverage begins to decrease. In practice, we usually evaluate the effectiveness of both methods on datasets, and incorporate our a priori knowledge on the hyper-parameter tuning phase. We may first find K for "Fixed K"(tune the first method) and when it comes to the finding (MIN_K, MAX_K), we can try to choose the $[MIN_K, MAX_K]$ interval in a way to contain the fixed K values found before.

$$MIN_K \leq \text{ fixed } K \leq MAX_K.$$

Once K is found, we try to decrease value of γ, which decreases the mean interval size. Our goal is to have the smallest mean tolerance interval size that satisfies our inclusion constraint. The idea is to fix the neighborhood parameters with the values found in the preceding process and decrease γ. This procedure is repeated as long as the inclusion constraint is satisfied. High values of γ will guarantee the satisfaction of the PIM tuning constraint but the computed intervals can be very large. Note that, with this approach, the value of γ can be less than β and this may happen when the local density of the response variable is quite dense. Based on the new value of γ, we can go to the first step and recalculate new values for the neighborhood hyper-parameter (K or the pair (MIN_K, MAX_K)) and this can be repeated for one or two iterations.

Algorithm 3 Hyper-parameter tuning for PIM with variable K.

function TuneHyper-Params($error\_set, \beta$)
  $\gamma \leftarrow 0.99$ ▷ or $\gamma \leftarrow 0.9$ depending on the dataset
  $(MIN_K, MAX_K) \leftarrow (MIN_{K_0}, MAX_{K_0})$ initial values
  $\lambda \leftarrow (\gamma, MIN_K, MAX_K)$
  for $iteration = 1..3$ **do**
    $(MIP, MIS) \leftarrow$ ComputeOnTrainigSet(β, λ)
    $MIS_{min} \leftarrow MIS$.
    while $MIP == \beta$ and $MIS \leq MIS_{min}$ **do**
      $(MIN_K, MAX_K) \leftarrow (MIN_K, MAX_K) + somestep$
      $\lambda \leftarrow (\gamma, MIN_K, MAX_K)$
      $MIS_{min} \leftarrow MIS$.
      $(MIP, MIS) \leftarrow$ComputeOnTrainigSet(β, λ)
    end while
    while $MIP == \beta$ and $MIS \leq MIS_{min}$ **do**
      $\gamma \leftarrow \gamma - step$
      $\lambda \leftarrow (\gamma, MIN_K, MAX_K)$
      $MIS_{min} \leftarrow MIS$.
      $(MIP, MIS) \leftarrow$ComputeOnTrainigSet(β, λ)
    end while
  end for
  return (MIN_K, MAX_K, γ)
end function

function ComputeOnTrainigSet(β, λ)
  Use Algorithms 1 and 2 to obtain tolerance intervals on the training set with a LOO or 10-fold cross-validation schema.
  $MIP \leftarrow$ use Equation (26) on the tolerance intervals found in the previous step.
  $MIS \leftarrow$ compute the mean size of tolerance intervals found above.
  return (MIP, MIS)
end function

8. Experiments

In this section we will use several regression datasets to compare our predictive interval method for local linear regression with other interval prediction methods. The selected methods will be tested upon their capacity to provide two-sided β-content predictive interval models. The models are compared for their reliability, efficiency, precision and the tightness of their obtained envelope as described in Section 4. Note that we are interested in comparing the mentioned methods regardless of any variable selection or outliers detection preprocessing. This section is organized in five parts: the first part describes our datasets, the second part describes the interval prediction methods that are used in the third part and the fourth part is a discussion of results.

8.1 Dataset description

In this work we use nine benchmark datasets to validate our suggested methods. These datasets are listed below, where we can find each dataset name in double quotes and its abbreviation in parentheses. Then we mention their numbers of predictor and number of instances, respectively denoted by p and n. Note that some of these datasets have fewer variables than their source because we systematically removed any instances having null values. The "Parkinsons Telemonitoring" dataset Frank and Asuncion (2010) contains two regression variables named "motor_UPDRS" and "total_UPDRS". We considered it as two distinct datasets named "Parkinson1" and "Parkinson2". Each dataset has one of the "motor_UPDRS" or "total_UPDRS" variables.

- "Parkinsons Telemonitoring" Frank and Asuncion (2010) (Parkinson1). We extracted Parkinson1 from the "Parkinsons Telemonitoring" Frank and Asuncion (2010) dataset. It has the "total_UPDRS" variable and does not contain the "motor_UPDRS" variable. $n = 5875, p = 21$.

- "Parkinsons Telemonitoring" Frank and Asuncion (2010) (Parkinson2). We extracted Parkinson2 from the "Parkinsons Telemonitoring" Frank and Asuncion (2010). It has the "motor_UPDRS" variable and does not contain the "total_UPDRS" variable. $n = 5875, p = 21$.

- "Wine Quality" Cortez et al. (1998) (Wine) (Red Wine). $n = 4898, p = 12$.

- "Concrete Compressive Strength" (Concrete) Yeh (1998). $n = 1030, p = 9$.

- "Housing" Frank and Asuncion (2010) (Housing). $n = 506, p = 14$.

- "Auto MPG" (Auto) Frank and Asuncion (2010). $n = 392, p = 8$.

- "CPU" Frank and Asuncion (2010) (CPU). $n = 209, p = 7$.

- "Concrete Slump Test" Yeh (2007) (Slump). $n = 103, p = 10$.

- "Motorcycle" (Motorcycle) Silverman (1985). $n = 133, p = 1$.

8.2 Interval Prediction Methods

In this section we describe the interval prediction methods used to build predictive interval models. Our experiments are performed with the R programming language. So, we first describe how each tested method is implemented in R. Each method has some general and dataset specific hyper-parameters. General hyper-parameter values are given next to the method name in the listing below, and the dataset specific hyper-parameters values are given in Table 1. Note that linear models do not have any hyper-parameters.

8.2.1 Method's Implementation

All the interval prediction methods listed below are explained in Section 3, except for our predictive-interval method for linear loess, which is introduced in Section 7. The selected methods are as follows:

- "Fixed K": two-sided predictive interval for linear loess as explained in 6 with the fixed K LHNPE neighborhood.
- "Var. K": two-sided predictive interval for linear loess as explained in 6 with the variable K LHNPE neighborhood.
- "LQR": two-sided interval prediction with linear quantile regression Koenker (2005). We used the *rq* and *rq.predict* function in R's *quanterg* package.
- "LQRC" two-sided Bonferroni 0.95-level confidence β-content interval obtained with two different quantile regression models as explained in "Confidence intervals on regression quantiles" of 3.4.2. We used the *rq* and *rq.predict* functions in R's *quanterg* package. We use *predict* with the following arguments: interval="confidence", type="percentile", se="boot", bsmethod= "wild".
- "NPQR": two-sided interval prediction by two non-parametric quantile regression models Takeuchi et al. (2006) as explained in "Quantile regression" of 3.4.2. This method's hyper-parameter minimizes the Pin-ball loss function with a 10-fold CV on the training set. This method is implemented by the *kqr* function in R's *kernlab* package. We use *kqr* with the following arguments: kernel="rbfdot", for a radial basis kernel function. We set kpar= "automatic" as the default value for radial basis functions. C=4, the cost regularization parameter is set between 3.8 and 5, depending on the dataset.
- "NPQR CV" : two-sided interval prediction by two non-parametric quantile regression models Takeuchi et al. (2006). The "NPQR CV" hyper-parameters are tuned in a way to find intervals that, in a 10-fold CV on the training set, have the smallest MIS and satisfy the tuning MIP constraint. We use the *kqr* function in R's *kernlab* package. We use *kqr* with the following arguments: kernel="rbfdot", for a radial basis kernel function. We set kpar= "automatic" as the default value for radial basis functions. C=0.1, the cost regularization parameter is chosen to lie 0.05 and 0.2, depending on the dataset. Satisfying the tuning MIP constraint on the training set requires us to select small values of cost regularization parameters.

- "LS-SVM Conv.": the conventional interval prediction method explained in 3.3 obtained with a least-square SVM regression. We used the *ksvm* function in *R*'s *kernlab* package. We use *ksvm* with the following arguments: kernel="rbfdot", for a radial basis kernel function. kpar= list(sigma= 0.2), the sigma hyper-parameter is set between 0.01 and 0.45, depending on the dataset, except for the motorcycle dataset which has sigma=6. We also set tau = 0.01, reduced = TRUE, tol = 0.0001.

- "Loess Conv." the conventional interval prediction method explained in 3.3 obtained with a linear loess regression.

We use the Tricube kernel, as in Cleveland and Devlin (1988), as the kernel function in all of our experiments.

8.2.2 Dataset Specific Hyper-Parameters

The linear Loess regression uses the K_{loess}-nearest neighbors as the bandwidth. This K_{loess} is found by minimizing the 10-fold cross validation error on the training set. For more details about linear loess see 2.3.3. All the non-linear methods listed above have at least one hyper-parameter that must be tuned on the dataset. These hyper-parameters are mentioned in Table 1 except for "Fixed K" and "Var. K", because "Fixed K" and "Var. K" may have different hyper-parameter for different β value. Thus their hyper-parameter values are mentioned with their method results.

Dataset	"NPQR" C	"NPQR CV" C	"LS-SVM Conv." sigma	"Loss Conv." K_{loess}
Parkinson1	5	0.2	0.25	80
Parkinson2	5	0.1	0.2	70
Wine	5	0.1	0.45	150
Concrete	4	0.1	0.3	80
Housing	4.5	1	0.08	60
Auto	3.8	0.2	0.25	30
CPU	4	0.2	0.025	40
Slump	4.5	0.05	0.05	30
Motorcycle	4	0.1	6	15

Table 1: Hyper-parameter values for non-linear interval prediction models.

8.2.3 Hyper-parameter tuning strategy

In a first attempt, datasets are divided into two subsamples of size $\frac{2}{3}n$ and $\frac{1}{3}n$, where n represents the dataset size. The part containing $\frac{2}{3}$ of observation is used to tune the predictive interval model's hyper-parameters. Then, all of the instances will serve to validate the results using a 10-cross validation scheme. **Note that we are interested in comparing the mentioned methods regardless of any variable selection or outlier detection preprocessing.**

8.3 Testing Predictive Interval Models

The goal of this section is to compare the above-mentioned interval prediction methods based on their strength while providing β-content predictive interval models. The models are compared based on **reliability, efficiency, precision and the tightness of their envelope**. Our introduced methods ("Var. K" and "Fixed K.") are used to obtain predictive interval models for Local Linear Regression (LLR). Consequently, we first compare our methods with the conventional interval prediction on the local linear regression ("Loess Conv."). For this purpose, we will use Tables 2 and 3 which compare "Loess Conv.", "Var. K" and "Fixed K.". These models are built upon the same regression model and their only difference is their interval computation algorithm. Our tables provide detailed experimental results but they take a lot of spaces which make them hard to interpret, and not useful for comparing several methods across different datasets. We will use MIP charts, MIS charts and EGSD charts to compare all of the interval prediction methods. This comparison measures a method's strength, while providing β-predictive interval models with $\beta = 0.8, 0.9, 0.95$ and 0.99. We have chosen five big datasets and compare in a very detailed manner the precision, reliability, efficiency and envelope width of our models with the conventional model which is its most efficient competitor.

8.3.1 Comparing Local linear Methods

Outliers, limited number of observations and contrast between our assumptions and the true regression function cause errors in the prediction process. These errors occur in a similar manner when estimating the response variable distribution and they increase with β. For $\beta = 0.9, 0.95$, and particularly for $\beta = 0.99$, it becomes a critical task to find an effective interval prediction procedure that is able to find an upper bound for inter-quantiles of $Y(x)$. However these inter-quantiles are the most used ones in machine-learning and statistical hypothesis-testing. Hence, we will compare the methods based on their strength, while providing β-predictive interval models with $\beta = 0.8, 0.9, 0.95$ and 0.99.

Tables 2 and 3 are used to display the direct dataset measures explained in Section 4, for each dataset. These tables compare models of "Loess Conv." , "Var K." and "Fixed K.". For each dataset, we have 12 models, (3 methods : "Loess Conv." , "Var K." and "Fixed K." $\times$ 4 β's value : $0.8, 0.9, 0.95$ and $0, 99$). These 12 models are built on the same regression model which is a linear Loess model with K_{loess} as its bandwidth. K_{loess} is represented next to the dataset's name and it is found by minimizing the 10-fold cross validation error on the training set. Then we will use charts to compare our local linear predictive interval models with the other methods.

8.3.2 Table description

In Tables 2 and 3, each combination of dataset and β has a cell which displays $F^{0.05}_{\beta,n}$ for the underlying experiment. We can see if a model satisfies its PIM test or not. If it does not satisfy this constraint a $\blacktriangleleft$ or $\lhd$ sign may appear. The $\blacktriangleleft$ sign appears when the current model is the only one to fail the PIM test. When more than one of the three compared model fails, the $\lhd$ sign is put near their results. For each experiment, the model which passes the PIM test and has the smallest MIS is distinguished with the * sign. If a method receives the * sign for two consecutive β of the same dataset, it is annotated in bold and

with * . When it comes to our introduced model hyper-parameters, "Var K." needs the value of MIN_K, MAX_K and γ and "Fixed K." has a proper value for its K and γ. These hyper-parameters are illustrated in each dataset row.

8.3.3 Table commentaries

By looking at Tables 2 and 3, one can see that the three methods work for $\beta = 0.8$ on benchmark dataset. When the desired proportion is 0.8, "Var K." is slightly more effective than "Loess Conv." and "Fixed K." finds the biggest intervals. When it comes to $\beta = 0.9$, "Loess Conv." loses its reliability and fails to satisfy the MIP constraint on three datasets. If we increase the desired proportion to 0.95, the situation stays the same for "Var K." and "Fixed K.", but "Loess Conv." becomes much more unreliable. In fact it fails to satisfy the MIP constraint for five of the nine datasets. When looking in more detail, one can observe that "Fixed K." has almost everywhere larger MIP and gives wider intervals than others. It is important to emphasize that **the conventional method is nowhere more reliable than our methods.** We can also observe that "Var K." usually appears with the * sign and it is the only method which becomes bold. It means that it usually works and obtains the tightest band.

8.4 Comparing All Methods by Charts

Our tables are not useful for displaying the eight methods listed in 8.2.1, so for the sake of readability we produced Figures 2, 3, 4 and 5. These figures are MIP charts for our experiments. We can see that our introduced methods obtain high MIP, but we need more information to compare their reliability and efficiency. For this purpose we will use the MIS ratio charts and EGSD charts that are explained in 4.3.

8.4.1 Chart description

Each β value has a MIP, an EGSD and a MIS ratio chart. For each β, its EGSD chart is displayed just after its MIS ratio chart. For example, Figure 6 is the MIS Ratio chart for $\beta = 0.8$ and just after Figure 7 is the EGSD chart for $\beta = 0.8$. The MIS ratio charts display the MIS ratio for the reliable models (models which pass the PIM test). For a given dataset, the method having the smallest MIS ratio value is that which finds the tightest reliable envelope (the set of all obtained intervals). The EGSD chart displays the normalized EGSD value for all models. For a given dataset, the model having the smallest EGSD value has an Equivalent Gaussian distribution with the smallest variance.

8.4.2 Chart commentaries

Now we can easily compare all the interval prediction methods. Figures 6 and 7 respectively display the MIS ratio chart and EGSD chart for $\beta = 0.8$. $\beta = 0.8$ is the easiest case and all the methods can provide reliable predictive interval model. One can observe that "Var. K" and "Fixed K" models are almost always more efficient than the others. If we look in more detail, we can see that "Var. K" usually finds both the smallest MIS ratio and EGSD value. The conventional methods "Loess Conv." and "LS-SVM Conv." are the next

Dataset	Method	80%			90%		
		MIP	MIS (σ_{is})	$F^{0.05}_{0.8,n}$	MIP	MIS (σ_{is})	$F^{0.05}_{0.9,n}$
Parkinson1 (n=5875, p=21), K_{loess}=80	Loess Conv.	86.99	4.96	79.14	90.08	6.37	89.35
	Fixed K = 40, $\gamma = 0.9$	91.55	5.48 (4.4)		94.88	7.04 (5.64)	
	Var. K *	88.55	4.39 (3.78) *		92.81	5.64 (4.85) *	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (20\ ,\ 60\ ,\ 0.9)$					
Parkinson2 (n=5875, p=21), K_{loess}=70	Loess Conv.	86.36	3.53	79.14	89.95	4.53	89.35
	Fixed K = 50, $\gamma = 0.9$	91.46	4.2 (3.22)		94.64	5.4 (4.14)	
	Var. K *	89.08	3.52 (2.95) *		93.03	4.52 (3.79) *	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (30\ ,\ 60\ ,\ 0.9)$					
Wine (n=4898, p=12), K_{loess}=150	Loess Conv.	80.45	1.58 *	79.05	88.39 ◀	2.03	89.29
	Fixed K = 50, $\gamma = 0.7$	82.88	1.75 (0.38)		90.62	2.25 (0.48) *	
	Var. K	83.19	1.77 (0.39)		91.09	2.27 (0.51)	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (30\ ,\ 60\ ,\ 0.9)$					
Concrete (n=1030, p=9), K_{loess}=80	Loess Conv.	79.89	16.63 *	77.94	87.56 ◀	21.35	88.46
	Fixed K = 35, $\gamma = 0.5$	82.61	16.76 (5.73)		91.45	21.52 (7.36) *	
	Var. K	83.68	17.03 (5.91)		93	22.2 (7.59)	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (20\ ,\ 60\ ,\ 0.9)$					
Housing (n=506, p=14), K_{loess}=60	Loess Conv.	87.17	8.47	76.67	92.68	10.88	87.5
	Fixed K = 40, $\gamma = 0.9$	87.97	8.67 (3.31)		92.7	11.14 (4.25)	
	Var. K *	84.59	7.8 (2.8) *		91.72	10.01 (3.6) *	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (30\ ,\ 55\ ,\ 0.9)$					
Auto (n=392, p=8), K_{loess}=30	Loess Conv.	89.29	8.56	77.07	93.61	10.98	87.8
	Fixed K = 50, $\gamma = 0.9$	88.27	7.76 (3.14)		94.41	9.96 (4.03)	
	Var. K *	84.7	6.91 (2.81) *		92.61	8.87 (3.61) *	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (30\ ,\ 60\ ,\ 0.9)$					
CPU (n=209, p=7), K_{loess}=40	Loess Conv.	79.87	84.64	75.44	85.13 ◀	108.63	86.58
	Fixed K = 40, $\gamma = 0.9$	85.16	88.07 (64.23)		91.4	113.04 (82.44)	
	Var. K *	80.37	78.49 (59.2) *		88.97	100.75 (20.89) *	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (20\ ,\ 50\ ,\ 0.9)$					
Slump (n=103, p=10), K_{loess}=30	Loess Conv.	87.63	5.58	73.51	91.45	7.17	85.13
	Fixed K = 20, $\gamma = 0.5$	85.72	4.85 (1.41)		88.54	6.23 (1.81)	
	Var. K *	83.81	4.32 (1.24) *		87.63	5.55 (1.6) *	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (15\ ,\ 30\ ,\ 0.5)$					
Motorcycle (n=133, p=1), K_{loess}=30	Loess Conv.	82.57	61.11	74.29	90.21	78.43	85.72
	Fixed K = 35, $\gamma = 0.7$	85.72	66.47 (32.87)		96.31	85.31 (42.2)	
	Var. K *	85.6	56.73 (25.27) *		94	72.82 (32.44) *	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (15\ ,\ 35\ ,\ 0.7)$					

Table 2: Predictive interval models for local linear regression built on benchmark datasets with $\beta = 0.9, \beta = 0.9$.

Dataset	Method	95%			99%		
		MIP	$MIS\ (\sigma_{is})$	$F^{0.05}_{0.95,n}$	MIP	$MIS\ (\sigma_{is})$	$F^{0.05}_{0.99,n}$
Parkinson1 (n=5875, p=21), K_{loess}=80	Loess Conv.	92.35 ◀	7.59	94.53	94.85 ◁	9.97	98.78
	Fixed K = 40, $\gamma = 0.99$	97.61	9.63 (7.7)		98.74 ◁	12.66 (10.2)	
	Var. K	96.31	7.72 (6.51) *		98.08 ◁	10.15 (8.56)	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (20\ ,\ 60\ ,\ 0.99)$					
Parkinson2 (n=5875, p=21), K_{loess}=70	Loess Conv.	91.91 ◀	5.4	94.53	94.41 ◁	7.09	98.78
	Fixed K = 50, $\gamma = 0.99$	97.4	7.26 (5.57)		98.64 ◁	9.54 (7.32)	
	Var. K	96.35	6.1 (5.07) *		98.13 ◁	8.02 (6.76)	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (30\ ,\ 60\ ,\ 0.99)$					
Wine (n=4898, p=12), K_{loess}=150	Loess Conv.	92.4 ◀	2.43	94.48	97.3 ◁	3.19	98.76
	Fixed K = 50, $\gamma = 0.9$	95.83	2.91 (0.63) *		98.75 ◁	3.82 (0.83)	
	Var. K	96.42	3.05 (0.67)		98.93	4.02 (0.88) *	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (30\ ,\ 60\ ,\ 0.9)$					
Concrete (n=1030, p=9), K_{loess}=80	Loess Conv.	93.87 ◀	25.44	93.88	98.53	33.43 *	98.49
	Fixed K = 35, $\gamma = 0.5$	95.62	25.64 (8.77) *		99.02	33.7 (11.53)	
	Var. K	95.72	26.46 (9.04)		99.02	34.77 (11.88)	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (20\ ,\ 60\ ,\ 0.9)$					
Housing (n=506, p=14), K_{loess}=60	Loess Conv.	95.24	12.96 *	93.18	97.62 ◀	17.04	98.17
	Fixed K = 40, $\gamma = 0.9$	95.45	13.27 (5.07)		98.61	17.44 (6.66) *	
	Var. K	96.24	13.8 (5.01)		98.61	18.14 (6.58)	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (30\ ,\ 50\ ,\ 0.99)$					
Auto (n=392, p=8), K_{loess}=30	Loess Conv.	96.17	13.09	93.4	97.46 ◀	17.2	98.27
	Fixed K = 50, $\gamma = 0.99$	97.2	13.39 (5.42)		98.71	17.6 (7.12)	
	Var. K *	96.44	11.99 (4.82) *		98.71	15.76 (6.34) *	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (30\ ,\ 60\ ,\ 0.99)$					
CPU (n=209, p=7), K_{loess}=40	Loess Conv.	86.11 ◀	129.45	92.52	91.39 ◁	170.12	97.86
	Fixed K = 40, $\gamma = 0.99$	96.16	154.67 (112.8)		98.07	203.27 (148.24) *	
	Var. K	94.25	137.68 (101.75) *		96.64 ◁	180.95 (133.72)	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (20\ ,\ 50\ ,\ 0.99)$					
Slump (n=103, p=10), K_{loess}=30	Loess Conv.	94.36	8.54	91.46	97.18 ◀	11.23	97.38
	Fixed K = 20, $\gamma = 0.9$	97.18	9.35 (2.72)		98.09	12.29 (3.57)	
	Var. K *	96.27	8.16 (2.25) *		98.09	10.73 (2.96) *	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (15\ ,\ 30\ ,\ 0.9)$					
Motorcycle (n=133, p=1), K_{loess}=30	Loess Conv.	93.23	93.46	91.89	98.51	122.82	97.58
	Fixed K = 35, $\gamma = 0.7$	97.8	101.66 (50.28)		99.23	133.6 (66.08)	
	Var. K *	96.31	86.77 (38.65) *		99.23	114.03 (50.8) *	
	Hyper. params.	$(MIN_K, MAX_K, \gamma) = (15\ ,\ 35\ ,\ 0.7)$					

Table 3: Predictive interval models for local linear regression built on benchmark datasets with $\beta = 0.95, \beta = 0.99$.

efficient ones. When it comes to testing $\beta = 0.9$, the situation stays almost the same for "Var. K" and "Fixed K" and "Var. K" remains the most efficient method. Conversely both the conventional methods fail to provide reliable predictive interval model for three of the nine datasets.

Figures 10 and 11 ($\beta = 0.95$) show that the conventional pair ("Loess Conv." and "LS-SVM Conv.") are definitely not reliable. Their non-working models find wider and less efficient envelope than our proposed models "Var. K" and "Fixed K". The scenario is still the same for "Var. K": It is the method which usually finds the tightest reliable envelope. It also provides models that, even compared to non-reliable models, have the smallest variance of prediction error. Finally let us look at Figures 12 and 13 ($\beta = 0.99$). In this case "Fixed K" and "LQRC" are the most reliable models. "Fixed K" takes second place. It fails once more than "Fixed K" and "LQRC". When comparing the efficiency, "Var. K" is still the most efficient solution but its gap decreases with others. In this case "Fixed K" becomes approximately as efficient as "Var. K". It is also interesting to note that, for $\beta = 0.99$, "LQRC" provides more efficient models than before.

Note also that both the conventional pair and the "NPQR CV" method fails more for large datasets than for small datasets. Small datasets do not have sufficient observations to reject the null hypothesis, which states that the tested model is a predictive interval model, so we accept their models as predictive interval models. It is also interesting to observe that "NPQR" fails in all cases. Its intervals are neither reliable nor efficient. These results are summarized in Table 4. Table 4 summarizes all the displayed charts. Each row of this table is dedicated to a different dataset which summarizes three qualities through $\beta = 0.8, 0.9, 0.95$ and $0, 99$. The first quality is the reliability: we cite the method which is the most reliable through $\beta = 0.8, 0.9, 0.95$. The second quality (the third column) shows the method that, for each dataset, generally provides the tightest reliable band and the fourth column displays the most efficient method. In the fourth column we ignore the method's reliability and we just compare its EGSD normalized value with others EGSD normalized value.

Dataset	Most Reliable	Tightest reliable band	General Efficiency (ignore the reliability)
Parkinson1	LQRC	Var. K	Var. K
Parkinson2	LQRC	Var. K	Var. K
Wine	Var. K & LQRC	Fixed K	Loess Conv
Concrete	Var K, Fixed K, LQRC, NPQR CV & Loess Conv	Fixed K.	Fixed K
Housing	Fixed K & LQRC	Fixed K	Fixed K
Auto	Var K, Fixed K, LQRC, NPQR CV & LS-SVM Conv	LS-SVM Conv	LS-SVM Conv
CPU	Fixed K	Var K	Var K
Slump	Var K, Fixed K, NPQR CV, Loess Conv & LS-SVM Conv	Var K	Var K
Motorcycle	Fixed K, Var K, NPQR CV, LS-SVM Conv & Loess Conv	Var K	Var K

Table 4: General ranking based on the MIP charts, MIS charts and EGSD charts for $\beta = 0.8, 0.9, 0.95$ and 0.99.

8.5 Detailed Comparison Using Plots

In the previous experiments, we concluded that our proposed predictive interval methods are the most reliable and effective method. Our goal is to compare in a very detailed manner the precision, reliability, efficiency and envelope tightness of our methods with one of its most efficient competitors. For this purpose, we have chosen the five largest datasets, because bigger datasets can provide more significant results. If we compare our introduced methods with the most reliable method, we have to select "LQRC". However "LQRC" is not more reliable than "Fixed K" but we have seen that "LQRC" is considerably less efficient and it only begins to be useful for $\beta > 0.95$. Therefore, we will have a more detailed comparison of our methods with the most effective interval prediction methods. We have seen that "Loess Conv." and then "LS-SVM Conv." are the most effective solutions after "Var K.". They have the same interval prediction methods but they use different regression algorithms, so we select "Loess Conv." which is revealed to be a bit more reliable and effective than "LS-SVM Conv." on the largest datasets. For this purpose we will use EGSD plots and MIP plots (described in 4.3) to compare "Var K.", "Fixed K." and "Loess Conv.".

8.5.1 Plot interpretation

For each dataset, the EGSD plot compares the efficiency of the tested models and the method having the highest line in this plot is the most inefficient one. MIS plots compares the envelope wideness of reliable models. The method having the most bottom line provides the most reliable envelope in the MIS plot. Note that in the MIS plot each model is plotted until its failure MIP. Once we have compared methods based on their envelope size and their efficiency, the MIP plot will help us to compare the precision and reliability of interval

prediction models. The best model in this plot is the one that has the nearest line to the upper side of the "Nominal MIP line". For further explanation of these plots, see 4.3.

8.5.2 Plot commentaries

By looking at Figure 14, we can see that when the nominal MIP is greater than or equal to 0.6, the "Loess Conv." model loses its efficiency but Figure 15 shows that the "Var K." model is the tightest reliable model. At the same time, Figure 16 compares their failure MIP. We can see that when the "Loess Conv." failure MIP is 0.93, our methods give a failure MIP equal to 0.99. This figure also shows that "Var K." has the most precise model and "Loess Conv." has the least reliable and precise one. The same experiment is performed for the Parkinson2 dataset which gives Figures 17, 18 and 19. In this case, "Loess Conv." becomes more efficient than "Fixed K." but it is always less efficient than "Var K.". Then Figures 18, shows that the "Var K." method provides again the most tightest reliable band. Next we look at the Parkinson2's MIP plot and we can see that the "Var K." model remains the most precise model. Figures 19 states that "Loess Conv." has a failure MIP of 0.93 and it is again the most unreliable solution.

In the Concrete dataset, the previous ranking gets better for "Fixed K.". "Loess Conv." is no longer the most efficient solution. Figures 20, 21 and 19 show that the "Fixed K." model is more effective and more precise than the "Loess Conv." model. The "Fixed K." method provides a model for the concrete dataset that obtains the tightest reliable band. Note that for $NominalMIP \geq 80$ ($\beta \geq 0.8$), its envelope is even tighter than the "Var K.". Our experiments continue with the Wine dataset where the "Loess Conv." model is the most efficient and provides the tightest band, however it has a failure MIP of 0.83 compared to a failure MIP of 0.99 for "Fixed K" and 0.97 for "Var K.".

We finalize our experiments with the Housing dataset where EGSD, MIS and MIP plots are displayed respectively in Figures 26, 27 and 28. Figures 26 shows that the "Var K." model, the "Loess Conv." model and the "Fixed K." model are respectively ranked as the first, the second and the third ranking efficient models. "Var K." is again the method that provides the tightest reliable band and for $\beta \geq 0.8$ the "Loess Conv." model is tighter than the "Fixed K." model. Figure 28 gives the same ranking for their precision. These rankings are summarized in Table 5. Each row of this table is dedicated to a different dataset which summarizes four qualities through 16 different inter-quantiles: $0.25 \leq \beta \leq 0.99$. The first three columns are similar to Table 4 except that they are obtained for $0.25 \leq \beta \leq 0.99$. The fourth column displays the method which is generally the most precise. This is the method that its MIP line, compared to other methods, remains the nearest to the upper side of the "Nominal MIP line".

Dataset	Most Reliable	Tightest reliable band	General Efficiency (ignore the reliability)	General Precision
Parkinson1	Var. K & Fixed K.	Var. K	Var. K	Var. K
Parkinson2	Var. K & Fixed K.	Var. K	Var. K	Var. K
Wine	Fixed K.	Loess Conv. for $\beta \leq 0.8$	Loess Conv.	Var. K
Concrete	Var. K & Fixed K.	Fixed K.	Var. K	Var. K.
Housing	Var. K & Fixed K.	Var. K	Var. K	Var. K

Table 5: General ranking based on the MIP plots, MIS plots and EGSD plots for $0.25 \leq \beta \leq 0.99$.

8.6 Discussion of Results

We have compared our introduced methods with six other well-known interval prediction methods. This comparison is performed with a 10-fold cross validation schema on nine benchmark regression datasets which contain between 1 and 21 predictors. For $\beta \geq 0.9$, it becomes a critical task to find an effective predictive interval method that works on all datasets. However these inter-quantiles are the most used ones in machine learning and statistical hypothesis testing. Hence, we first compared the mentioned methods based on their strengths while providing β-predictive interval models with $\beta = 0.8, 0.9, 0.95$ and 0.99. While comparing our methods with their six competitors, we found them to be the most reliable non-linear predictive interval models. Our experiments have shown that they are usually also the most effective solution.

The conventional methods "Loess Conv." and "LS-SVM Conv.',' as shown theoretically in Appendix B, are revealed to be unreliable solutions. They even fail for $\beta = 0.9$ although they are almost always less efficient than "Var K." and "Fixed K." and **their envelope is almost always larger than the "Var K." model's band**. There is just one case where "Fixed K." and "LQRC" are more reliable than "Var K.". However "LQRC" always provides much wider bands than our methods and it is also much more inefficient and imprecise than "Var K." and "Fixed K.". On the other hand, if we ignore their reliability, "Loess Conv." and "LS-SVM Conv." rank are the most efficient methods after "Var K.". They sometimes provide tighter bands than "Fixed K." , however a model which provides a tight band but usually does not work is not appropriate for predictive interval models. "NPQR CV" is more reliable than the conventional pair but it is the least efficient solution. "NPQR" and "LQR" are absolutely not appropriate for high confidence interval prediction.

In a second attempt, we compared our methods with their most effective competitor. These comparisons have been performed on the five largest datasets of the nine benchmarks and each time on 16 distinct desired contents (β value). These experiments show the superiority of "Var K." and then "Fixed K.". **We have seen that "Var K." usually provides models with the tightest bands and they are almost always the most effective and more precise than others.** "Fixed K." models are usually more effective

and precise than "Loess Conv.". Note that for $\beta \geq 0.5$, "Loess Conv." is the most effective solution but it is in the same time the least precise. By more effective, we mean that the normalized EGSD value is the smallest for $\beta \geq 0.5$ but does not provide the tightest and the most precise envelope. Thus we do not recommend "Loess Conv.", because its model provides intervals that are too wide.

In a regression context, the conditional mean, the conditional variance and/or the conditional quantile may have different functions. The conditional mean is the general trend of the regression function whereas the conditional quantile is more related to the local distribution of the response variable. Least-squares based interval prediction methods ("Loess Conv.", "LS-SVM Conv.", "Fixed K." and "Var. K") try to indirectly estimate the conditional quantile function. They first estimate the conditional mean and then, based on this estimated conditional mean, they estimate the conditional quantile. On the other hand, we have quantile regression based methods ("NPQR", "NPQR CV", "LQRC" and "LQR") which directly estimate the conditional quantile. We know that the general trend is easier to predict and its estimator, compared to the conditional quantile, has a higher speed of convergence. This is why all of our least-squares based interval prediction methods are more efficient than the quantile regression based methods. Another reason for this superiority may be the absence of a global conditional quantile function. It can occur where the conditional variance of the error distribution is not a global function of the predictors. Our proposed methods are in the class of least-squares based interval prediction methods, so they take advantage of this fast convergence. However they are more reliable and efficient than the other member of this class (conventional methods). This is because our methods take into account the sample size and find confidence intervals on inter-quantile of the local distribution for the response variable whereas the conventional methods just estimate asymptotic global inter-quantiles of the conditional response variable.

9. Discussion and Conclusion

This paper introduced predictive interval models for non-parametric regression. These models provide intervals which contain at least a desired proportion of the conditional distribution of the response variable given specified combination of predictors. They can be obtained with tolerance intervals for regression or confidence intervals for regression quantiles but the application of these methods in the non-linear and particularly the non-parametric case are limited in the literature. The originality of this work is to extend this concept to local linear models. Our method does not neglect the regression bias and finds intervals that work properly with biased regression models. We applied our method for two-sided interval predictions but they can also be used in a one-sided interval prediction context. We also introduced a statistical Predictive Interval Model (PIM) test. In the same context, we introduced two measures for ranking interval prediction models. These measures rate the efficiency and the tightness of the obtained envelope. In the experiments we compared our proposed predictive interval with 6 other interval prediction methods. The results show that our approach performs very well. Our methods rank usually better than other methods and remain the most reliable non-linear interval prediction methods in the experiments. We also provided the state of the art on the interval prediction methods

in regression. Figure 9 displays the positions of our methods compared to the state of the art.

Our predictive intervals are based on local linear regression. We assume that the mean regression function is locally linear and the prediction error variable has locally the same distribution. The idea behind this method is to exploit the local density of prediction error in the LHNPE neighborhood of the query point to find the most appropriate intervals that contain the desired proportion of response values. For this purpose, we use tolerance intervals on prediction errors and they are obtained with a fixed and variable neighborhood method. The prediction errors are obtained based on a local linear estimation which could be done by any regression bandwidth selection technique. Once the mentioned errors have been found, we can use them to obtain our non-parametric predictive intervals. For this purpose, we need a second bandwidth, which is the tolerance interval bandwidth (LHNPE bandwidth). The LHNPE bandwidth is generally included in the regression bandwidth.

Our method differs from conventional least-squares approaches for finding confidence intervals on the unknown conditional response variable because our method takes into account the sample size and finds confidence intervals on inter-quantile of the local distribution for the conditional response variable, while the conventional methods just estimate asymptotic global inter-quantiles for the conditional response variable.

Contrary to quantile regression, our method is based on the local linear least squares model, so one can obtain both the conditional mean function and the predictive intervals. Another main difference is that quantile regression obtains estimates of quantiles which, on average, estimate the true quantile function but our method proposes predictive intervals which contain at least a desired proportion of the conditional response variable. Quantile regression may sometimes be more robust than least-squares estimators but it suffers from several problems. One of these problems is the absence of a conditional quantile function. It can occur where the conditional variance of the error distribution is not a function of predictors. Now consider the case when the conditional quantile function is different to the conditional mean function. We know that the conditional mean estimator converges faster than the conditional quantile estimator Koenker (2005). Thus estimating intervals by quantile regression may be less efficient than using least-squares methods. Besides, it is important to note that quantile regression also suffers from the crossing quantile problem, which is not present here. Our proposed methods are in the class of least-squares based interval prediction methods, so they take advantage of their fast convergence. However they are more reliable and efficient than the other members of this class (conventional methods). This is because our methods take into account the sample size and find confidence intervals on inter-quantiles of the local distribution for the response variable whereas the conventional methods just estimate asymptotic global inter-quantiles of the conditional response variable.

General remarks

The advantages, drawbacks and limitations of our methods are listed below:

Advantages

- It is a reliable interval prediction method for local linear least squares models.

- It does not ignore the non-parametric regression bias.
- It can be used with models having heteroscedastic errors.
- It directly addresses the problem of having predictive intervals that contain at least the desired proportion of response values. It is not designed to work asymptotically and also works with small datasets.
- It does not suffer from the crossing quantiles effect.
- It provides one model for two-sided interval prediction.
- It is simple, reliable and effective.
- It is based on local linear regression, which is a well-known regression method.

Drawbacks

- It is limited to local linear regression.
- It has a greater computational complexity than conventional and quantile regression interval prediction methods.

Limit of Applications

In the following cases, our method may have **similar results to its alternatives**:

- For interval prediction models which contain a very high proportion (0.99 or more) of the distribution of $Y(x)$.
- The dataset is almost identically distributed in the feature space.
- The dataset has quite low heteroscedasticity.

Our methods are not suited when:

- The reliability of the predicted intervals is not a concern.
- There exits regression models having significantly better prediction results than non-parametric regression models.
- The prediction errors are not normally distributed.

For future work, The most promising idea is the extension of our predictive intervals to predictive interval on any regression function, i.e. support vector machines. Another horizon may be its generalization to the one-sided interval prediction problem. One can also apply our methods to interval prediction in time series models.

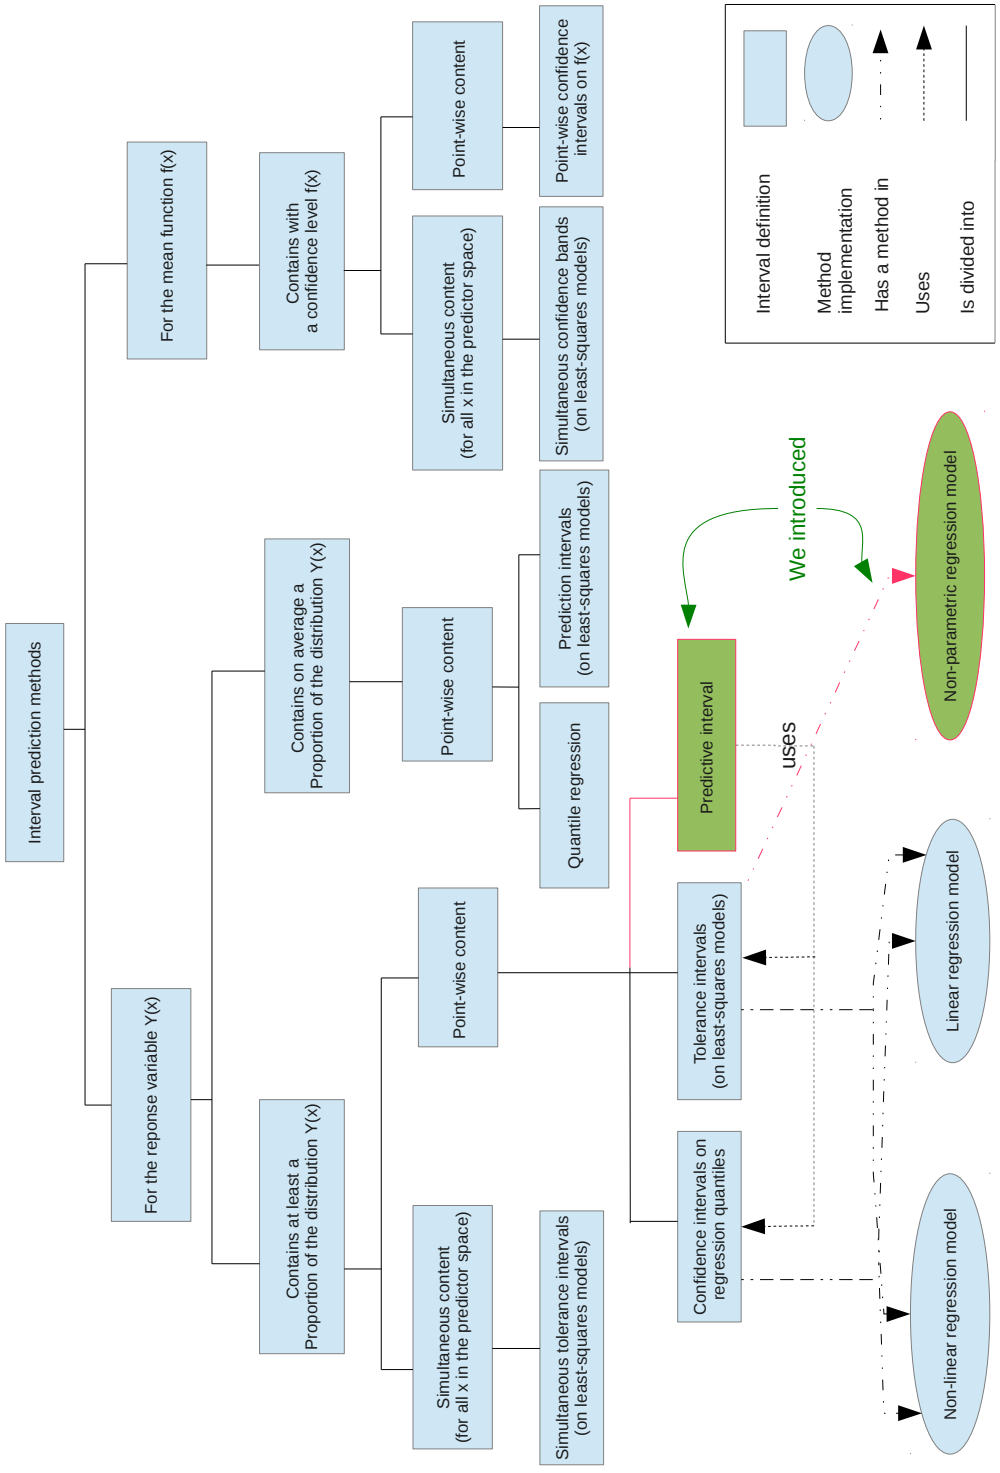
Interval prediction methods
For the mean function f(x)
Contains with a confidence level f(x)
Point-wise content
Point-wise confidence intervals on f(x)
Simultaneous content (for all x in the predictor space)
Simultaneous confidence bands (on least-squares models)
For the reponse variable Y(x)
Contains on average a Proportion of the distribution Y(x)
Point-wise content
Prediction intervals (on least-squares models)
Quantile regression
Contains at least a Proportion of the distribution Y(x)
Point-wise content
Simultaneous content (for all x in the predictor space)
Simultaneous tolerance intervals (on least-squares models)
Predictive interval
We introduced
uses
Non-parametric regression model
Tolerance intervals (on least-squares models)
Linear regression model
Confidence intervals on regression quantiles
Non-linear regression model
Interval definition
Method implementation
Has a method in
Uses
Is divided into

Figure 1: The position of the predictive intervals in the state of the art.

References

C. G. Atkeson, A. W. Moore, and Schaal S. Locally weighted learning. *Artificial Intelligence Review*, pages 11–73, 1997.

Raymond J Carroll and David Ruppert. Prediction and tolerance intervals with transformation and/or weighting. *Technometrics*, 33(2):197–210, 1991.

W. S. Cleveland. Robust locally weighted regression and smoothing scatterplots. *Journal of the American Statistical Association*, 74(368):829–836, 1979. ISSN 01621459.

W. S. Cleveland and S. J. Devlin. Locally weighted regression: An approach to regression analysis by local fitting. *Journal of the American Statistical Association*, 83(403):596–610, 1988. ISSN 01621459.

P. Cortez, A. Cerdeira, F. Almeida, T. Matos, and J. Reis. Modeling wine preferences by data mining from physicochemical properties. *Decision Support Systems*, 47(4):547–553, 1998.

J. Fan and I. Gijbels. *Local Polynomial Modelling and Its Applications: Monographs on Statistics and Applied Probability 66*. Monographs on Statistics and Applied Probability, 66. Chapman & Hall, 1996. ISBN 9780412983214.

J. Fan, N. E. Heckman, and M. P. Wand. Local polynomial kernel regression for generalized linear models and quasi-likelihood functions. *Journal of the American Statistical Association,90*, pages 141–150, 1995.

A. Frank and A. Asuncion. UCI machine learning repository, 2010. URL `http://archive.ics.uci.edu/ml`.

M. Ghasemi Hamed, M. Serrurier, and D. Durand. Simultaneous interval regression for k-nearest neighbor. In *Australasian Conference on Artificial Intelligence*, pages 602–613, 2012.

G. J. Hahn and W. Q. Meeker. *Statistical Intervals: A Guide for Practitioners*. John Wiley and Sons, 1991.

W. Härdle. *Applied nonparametric regression*. Econometric Society Monographs (No. 19). Cambridge University Press, 1990.

W. G. Howe. Two-sided tolerance limits for normal populations, some improvements. *Journal of the American Statistical Association*, 64(326):610–620, 1969. ISSN 01621459.

M. Kocherginsky, X. He, and Y. Mu. Practical confidence intervals for regression quantiles. *Journal of Computational and Graphical Statistics*, 14(1):41–55, 2005. ISSN 10618600.

R. Koenker. Confidence intervals for regression quantiles. In *Asymptotic statistics*, pages 349–359. Springer, 1994.

R. Koenker. *Quantile Regression*. Econometric Society Monographs. Cambridge University Press, 2005. ISBN 9780521608275.

K. Krishnamoorthy and T. Mathew. *Statistical Tolerance Regions: Theory, Applications, and Computation*. Wiley Series in Probability and Statistics. Wiley, 2009. ISBN 9780470473894.

E. Paulson. A note on tolerance limits. *The Annals of Mathematical Statistics*, 14(1):90–93, 1943. ISSN 00034851.

B. W. Silverman. Some aspects of the spline smoothing approach to non-parametric regression curve fitting. *Journal of the Royal Statistical Society. Series B (Methodological)*, 47 (1):1–52, 1985. ISSN 00359246.

C. J. Stone. Consistent nonparametric regression. *Annals of Statistics*, 5(4):595–620, 1977. ISSN 00905364.

I. Takeuchi, Q. V. Le, T. D. Sears, and A. J. Smola. Nonparametric quantile estimation. *Journal of Machine Learning Research*, 7:1231–1264, December 2006. ISSN 1532-4435.

P. Walley. *Statistical reasoning with imprecise probabilities*. Chapman and Hall, 1991.

W. A. Wallis. Tolerance intervals for linear regression. In *Proceedings Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 43–51. Univ. of Calif. Press, 1951.

I.-C. Yeh. Modeling of strength of high-performance concrete using artificial neural networks. *Cement and Concrete Research*, 28(12):1797–1808, 1998. ISSN 0008-8846.

I-Cheng Yeh. Modeling slump flow of concrete using second-order regressions and artificial neural networks. *Cement and Concrete Composites*, 29(6):474–480, 2007. ISSN 0958-9465.

Appendix A.

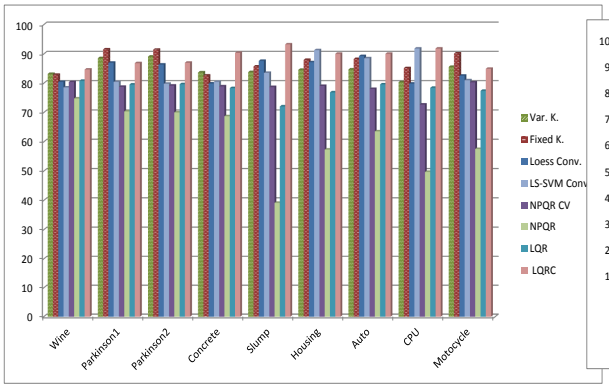

Figure 2: MIP chart for benchmark datasets with $\beta = 0.8$.

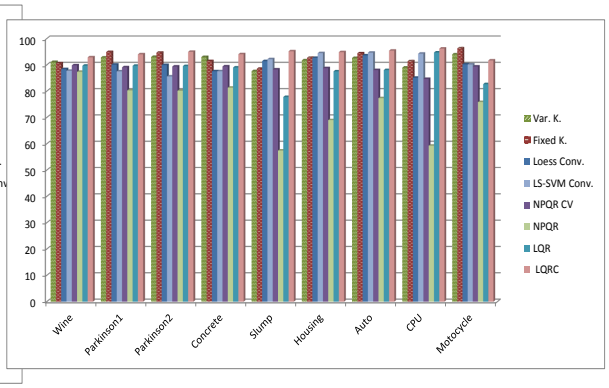

Figure 3: MIP chart for benchmark datasets with $\beta = 0.9$.

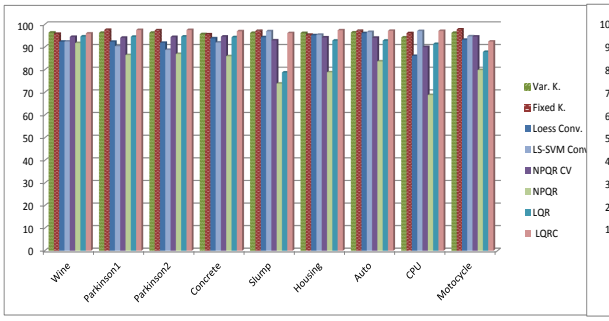

Figure 4: MIP chart for benchmark datasets with $\beta = 0.95$.

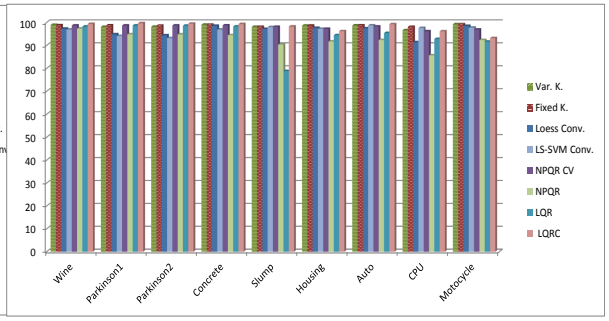

Figure 5: MIP chart for benchmark datasets with $\beta = 0.99$.

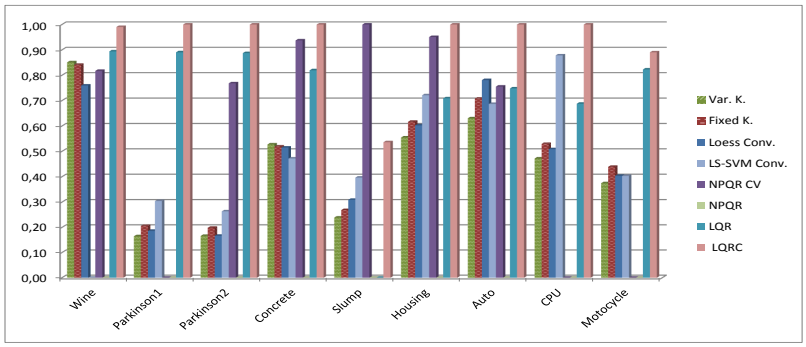

Figure 6: MIS Ratio chart for benchmark datasets with $\beta = 0.8$. The smallest value denotes the tightest reliable band.

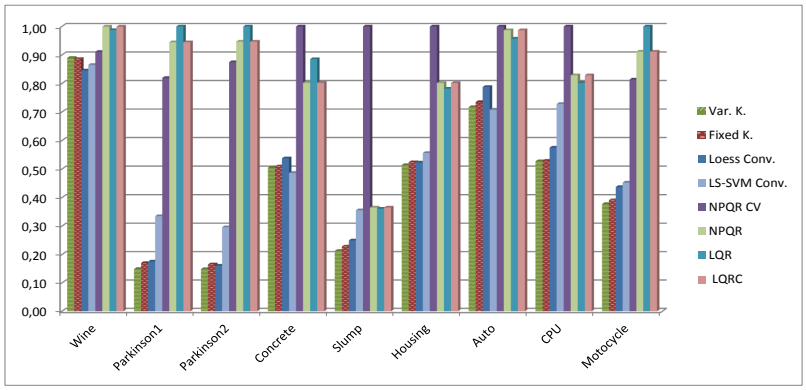

Figure 7: EGSD chart for benchmark datasets with $\beta = 0.8$. The smallest value denotes the most efficient band. This measure ignores the reliability.

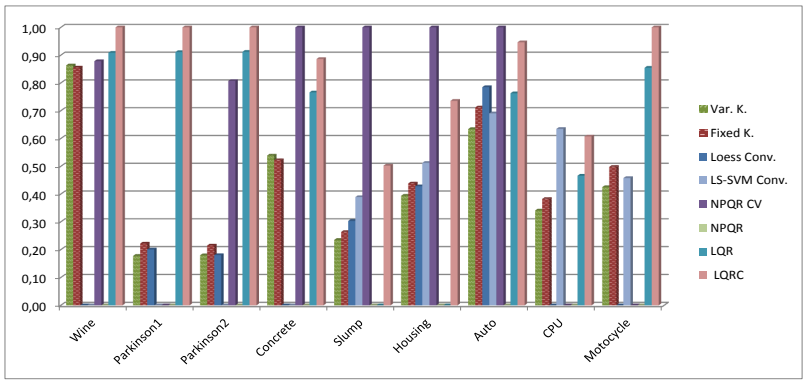

Figure 8: MIS Ratio chart for benchmark datasets with $\beta = 0.9$. The smallest value denotes the tightest reliable band.

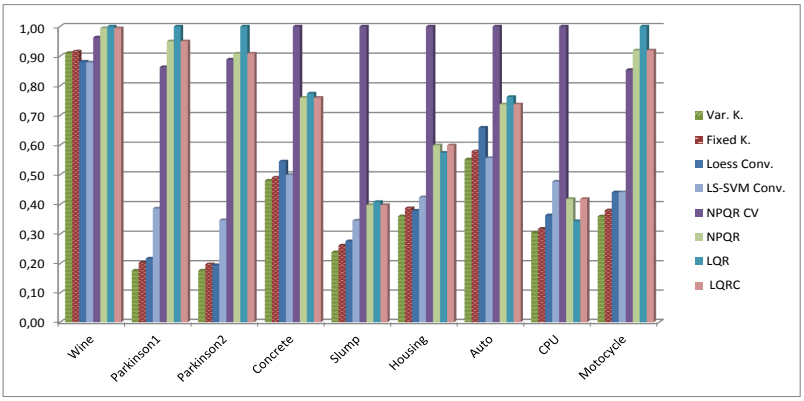

Figure 9: EGSD chart for benchmark datasets with $\beta = 0.9$. The smallest value denotes the most efficient band. This measure ignores the reliability.

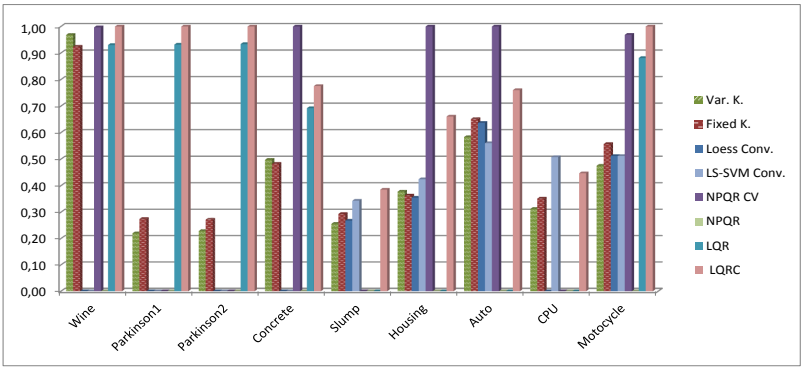

Figure 10: MIS Ratio chart for benchmark datasets with $\beta = 0.95$. The smallest value denotes the tightest reliable band.

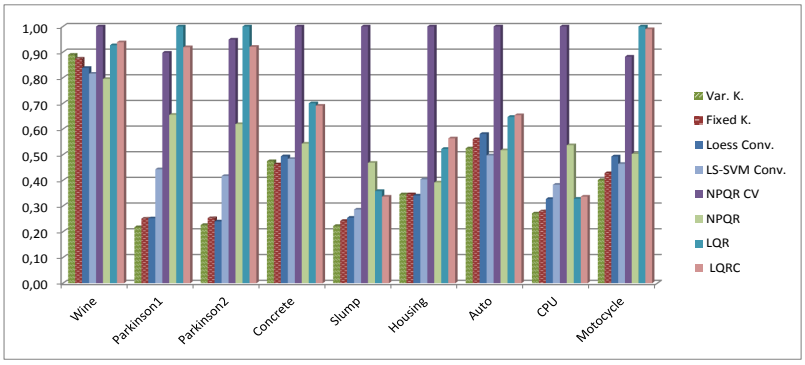

Figure 11: EGSD chart for benchmark datasets with $\beta = 0.95$. The smallest value denotes the most efficient band. This measure ignores the reliability.

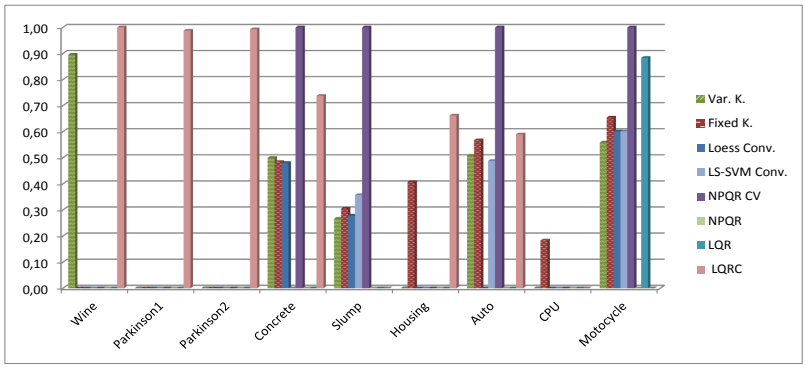

Figure 12: MIS Ratio chart for benchmark datasets with $\beta = 0.99$. The smallest value denotes the tightest reliable band.

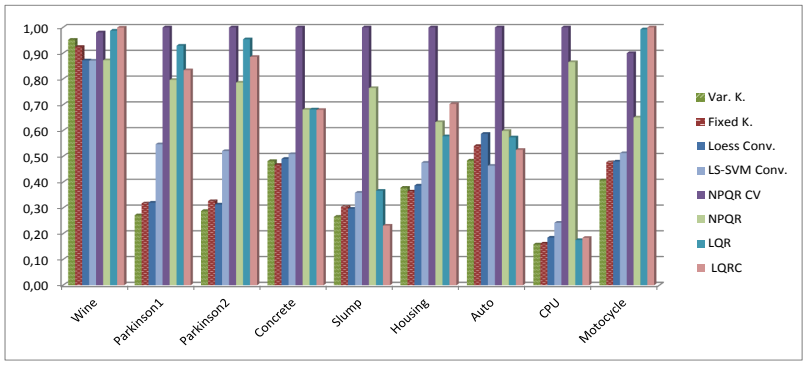

Figure 13: EGSD chart for benchmark datasets with $\beta = 0.99$. The smallest value denotes the most efficient band. This measure ignores the reliability.

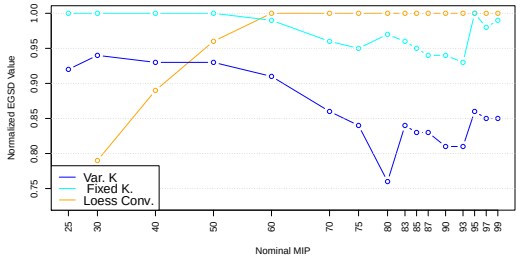

Figure 14: EGSD plot for Parkinson1 datasets. The lowest line denotes method that yields most efficient band. This measure ignores the reliability.

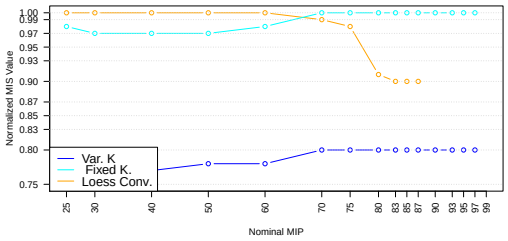

Figure 15: MIS plot for Parkinson1 datasets. The smallest value denotes the tightest reliable band.

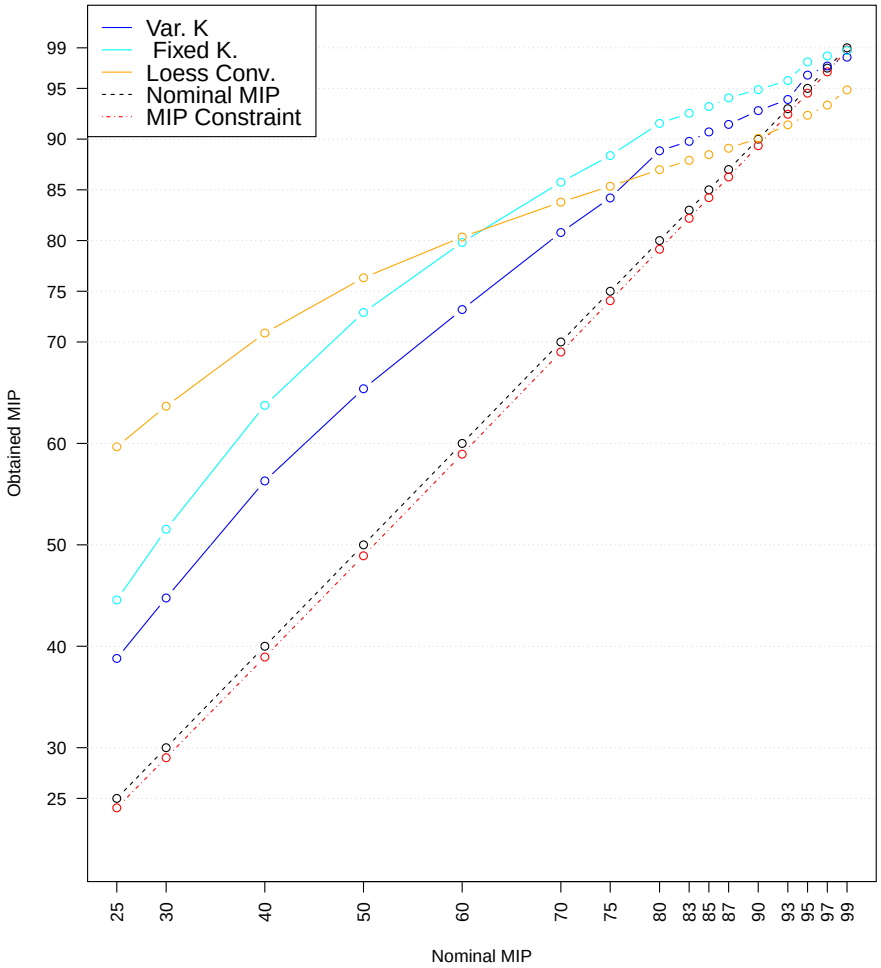

Figure 16: MIP plot for Parkinson1 datasets.

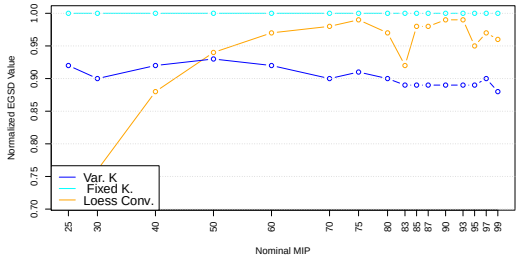

Figure 17: EGSD plot for Parkinson2 datasets. The lowest line denotes method that yields most efficient band. This measure ignores the reliability.

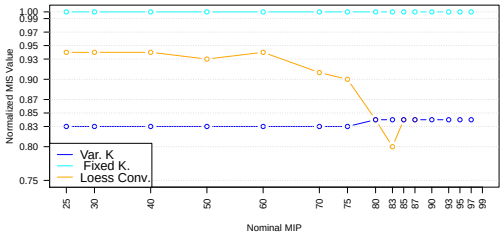

Figure 18: MIS plot for Parkinson2 datasets. The smallest value denotes the tightest reliable band.

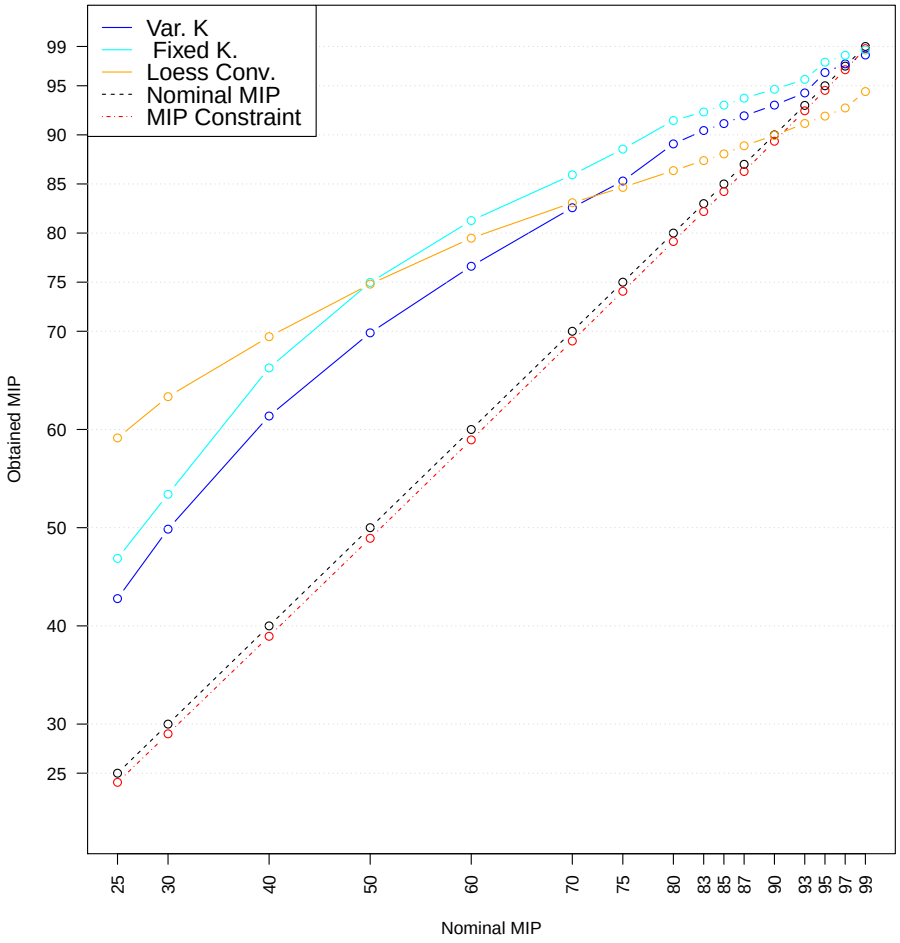

Figure 19: MIP plot for Parkinson2 datasets.

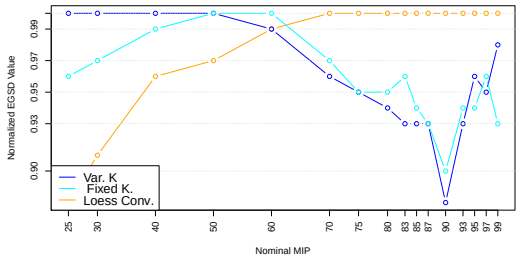

Figure 20: EGSD plot for Concrete datasets. The lowest line denotes method that yields most efficient band. This measure ignores the reliability.

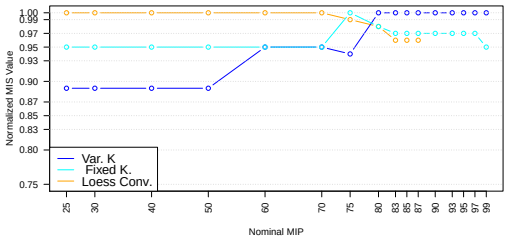

Figure 21: MIS plot for Concrete datasets. The smallest value denotes the tightest reliable band.

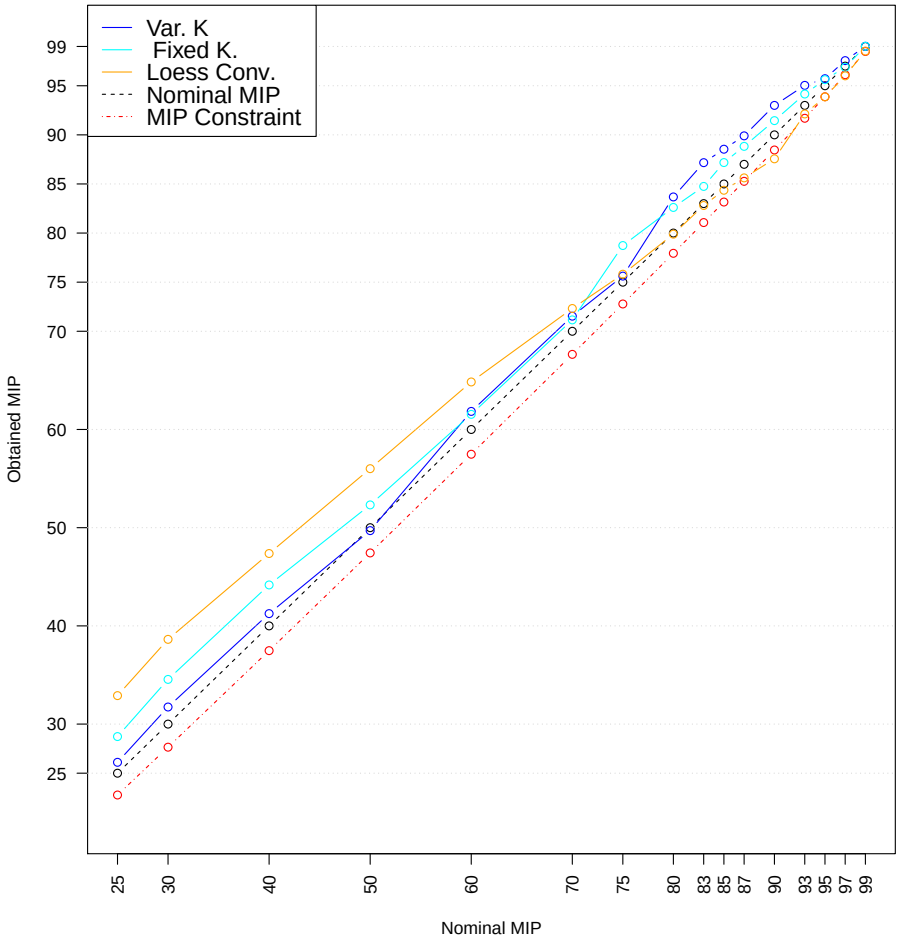

Figure 22: MIP plot for Concrete datasets.

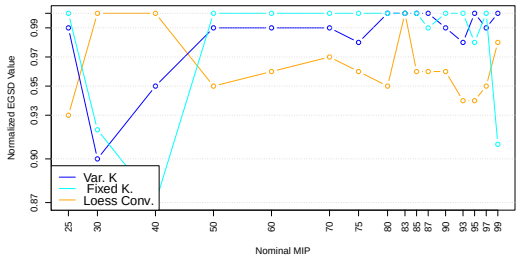

Figure 23: EGSD plot for Wine datasets. The lowest line denotes method that yields most efficient band. This measure ignores the reliability.

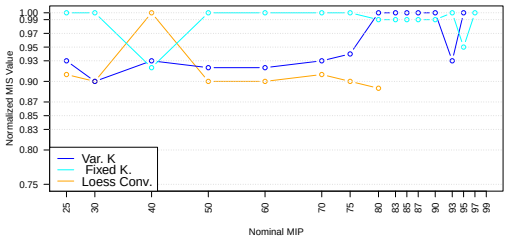

Figure 24: MIS plot for Wine datasets. The smallest value denotes the tightest reliable band.

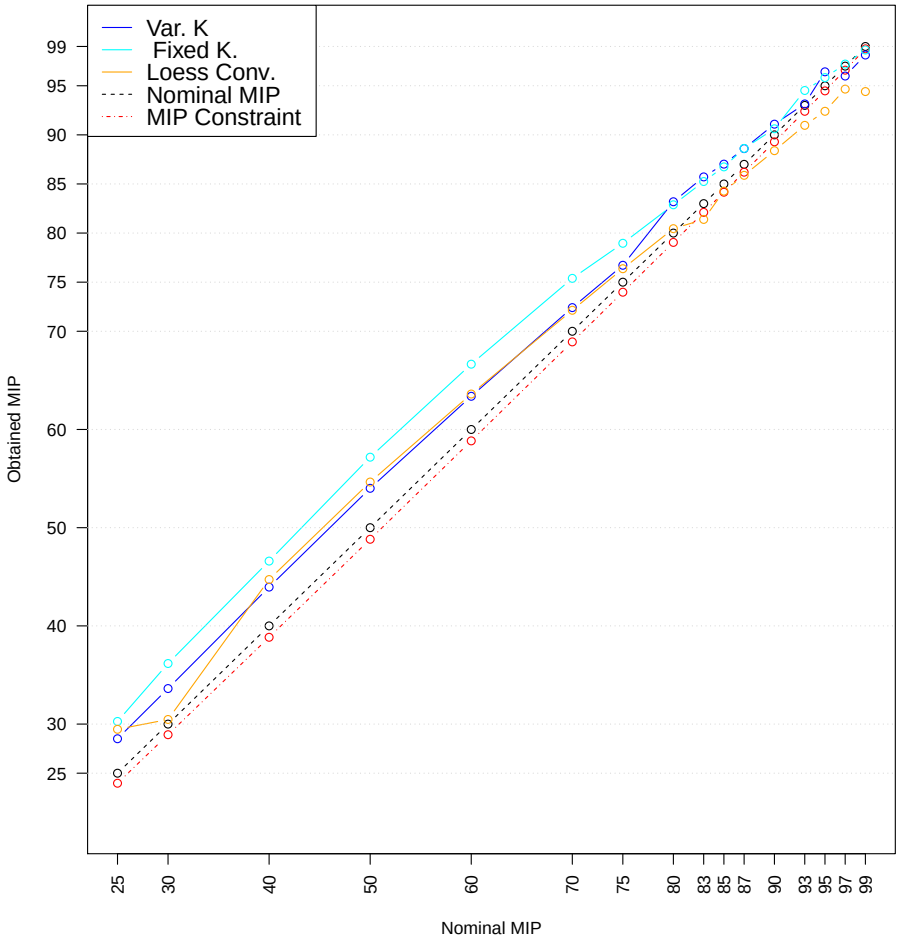

Figure 25: MIP plot for Wine datasets.

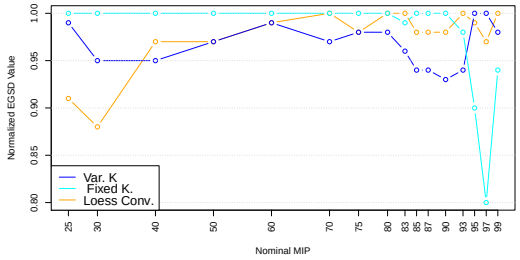

Figure 26: EGSD plot for Housing datasets. The lowest line denotes method that yields most efficient band. This measure ignores the reliability.

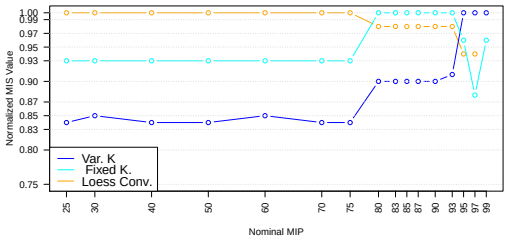

Figure 27: MIS plot for Housing datasets. The smallest value denotes the tightest reliable band.

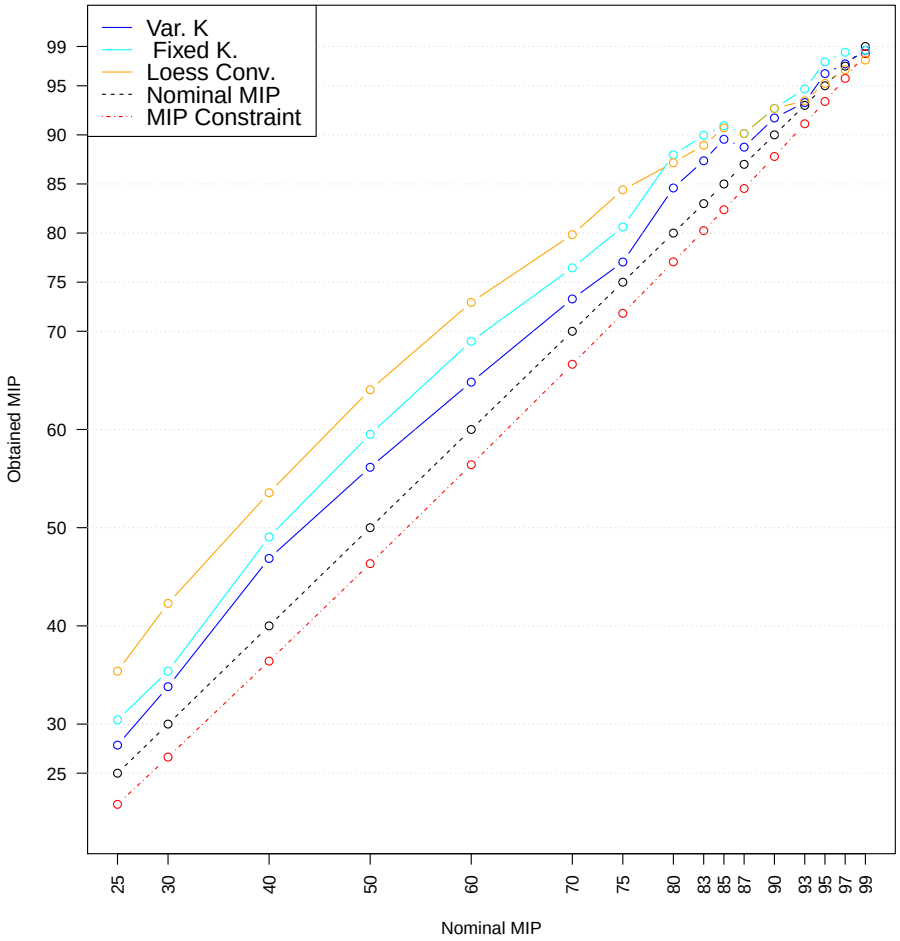

Figure 28: MIP plot for Housing datasets.

Appendix B.

.1 Predictive risk

Predictive risk is one of the most commonly used measure for tuning of hyper-parameters, model selection and inference. The risk of an estimator is the square of the difference between the true value of the parameter and its estimation. Given a fixed value of x, the mean of squared error for all values of the random variable $\hat{f}(x)$ is defined as the risk at point x:

$$MSE_{\hat{f}(x)} = RISK_{\hat{f}(x)} = E[(f(x) - \hat{f}(x))^2]. \tag{42}$$

It is well known that (42) can also be decomposed in bias and variance terms as in (43).

$$MSE_{\hat{f}(x)} = E[(\hat{f}(x) - f(x))^2] = Bias^2_{\hat{f}(x)} + \sigma^2_{\hat{f}(x)} \tag{43}$$

$$\text{, where } Bias_{\hat{f}(x)} = E[\hat{f}(x) - f(x)] \tag{44}$$

$$\text{and } \sigma^2_{\hat{f}(x)} = E[\hat{f}(x)^2] - E[\hat{f}(x)]^2. \tag{45}$$

Average Mean Square Error of $\hat{f}(\cdot)$, or the average risk of $\hat{f}(\cdot)$, is the average of the mean squared error of $\hat{f}(\cdot)$ over all values of x, and it is used as an evaluation measure in regression problems.

$$\text{Average MSE } = \frac{1}{n}\sum_{i=1}^{n} E[(\hat{f}(x_i) - f(x_i))^2].$$

The average risk is related to the predictive risk. Let us first define the squared prediction error. The squared prediction error is the squared error of prediction for a new observation (x_i, Y^*) and it is defined as:

$$(Y^* - \hat{f}(x_i))^2 = (f(x_i) + \epsilon^* - \hat{f}(x_i))^2.$$

The predictive risk is:

$$\begin{aligned}\text{Predictive Risk} &= \frac{1}{n}\sum_{i=1}^{n} E[(Y^* - \hat{f}(x_i))^2] \\ &= \text{Average MSE } + \frac{1}{n}\sum_{i=1}^{n} \sigma^2(x_i),\end{aligned}$$

so we have:

$$\text{Predictive Risk} = \text{Average MSE } + c, \text{ where } c = \frac{1}{n}\sum_{i=1}^{n} \sigma^2(x_i); \tag{46}$$

$\sigma^2(x_i)$ is the variance of the response variable at x_i, and c is a constant. If the error variance $\sigma^2(x_i)$ is constant for all x_i, then

$$\text{Predictive Risk} = \text{Average MSE } + \sigma^2. \tag{47}$$

Hence based on (46), minimizing the predictive risk results in minimizing the average risk of the estimated regression function $\hat{f}(\cdot)$. *In a small to medium size dataset, leave-one-out or 10-fold cross validation MSE are well-known estimators of predictive risk.*

.2 Conventional Interval prediction

One of the most common interval prediction techniques used in practice is to take $[\hat{f}(x) - Z_{\frac{1-\beta}{2}} SSE, \hat{f}(x) + Z_{1-\frac{1-\beta}{2}} SSE])$ as the interval which contains a β proportion of $Y(x)$'s population,where SSE^2 is the average MSE given by a Leave-One-Out (LOO) or a 10-fold cross validation scheme. One might assume that the intervals expressed below have similar properties to the regression tolerance interval defined in the next section.

$$P\left(Y(x) \in \left[\hat{f}(x) - Z_{\frac{1-\beta}{2}} SSE, \hat{f}(x) + Z_{1-\frac{1-\beta}{2}} SSE\right]\right) = \beta. \tag{48}$$

As seen in (46), the expected value of S^2 is approximately equal to the predictive risk. Thus, based on the bias-variance decomposition:

$$\begin{aligned} E(SSE^2) &= \text{Average MSE } + \frac{1}{n}\sum_{i=1}^{n} \sigma^2(x_i) \\ &= \frac{1}{n}\sum_{i=1}^{n} \text{Bias}^2_{\hat{f}(x_i)} + \frac{1}{n}\sum_{i=1}^{n} \sigma^2_{\hat{f}(x_i)} + \frac{1}{n}\sum_{i=1}^{n} \sigma^2(x_i) \end{aligned}$$

and we have:

$$E(SSE^2) = \text{Average Bias}^2_{\hat{f}(x_i)} + \text{Average } \sigma^2_{\hat{f}(x_i)} + \text{Average } \sigma^2(x_i) \tag{49}$$

We assume that:

- the error variance for all x's is constant (homoscedasticity).
- the estimator's variance $\sigma^2_{\hat{f}(x)}$ is constant for all x.
- $\hat{f}(x)$ is an unbiased estimator of $f(x)$.
- the error ε, and $\hat{f}(x)$, are independent and both have normal distributions.

Then, we have:

$$\begin{aligned} E(SSE^2) &= \sigma^2_{\hat{f}(x)} + \sigma^2 \\ \hat{f}(x) - \varepsilon &\sim \mathcal{N}(f(x), \sigma^2_{\hat{f}(x)} + \sigma^2). \end{aligned}$$

Considering the fact that n tends to be large and under the above condition, we can consider SSE^2 as an approximation to the variance of the prediction around any point $SSE^2 \approx \sigma^2_{\hat{f}(x)} + \sigma^2$, which results in (50):

$$\frac{\hat{f}(x) - Y(x)}{SSE} \sim \mathcal{N}(0, 1). \tag{50}$$

Thus under the above conditions, (48) becomes asymptotically valid, but it remains non-applicable for finite sample size datasets. The problem is addressed by tolerance intervals in least-squares models.

.3 Two-sided Bonferroni Confidence intervals on regression quantiles

These two-sided intervals contain with a γ confidence level, a proportion $1-\alpha$ of $Y(x)$. As noted, we need a pair of $(\frac{\alpha}{2}, 1-\frac{\alpha}{2})$ quantile regression models but each model now itself needs a one-sided confidence interval on regression quantile. Once we have built the upper and lower quantile regression model, we must obtain a lower (one-sided) $(1-\frac{\tau}{2})$ confidence interval on the lower $\frac{\alpha}{2}$-quantile regression model and an upper (one-sided) $(1-\frac{\tau}{2})$ confidence interval on the upper $(1-\frac{\alpha}{2})$-quantile regression model. By applying the Bonferroni inequality, one can merge the pair of $(1-\frac{\tau}{2})$ confidence intervals to obtain a joint confidence statement with a probability greater or equal to $\gamma = 1-\tau$. The lower and upper $(1-\frac{\tau}{2})$-confidence intervals are respectively denoted $IL_{\frac{\alpha}{2}}^{1-\frac{\tau}{2}}(x)$ and $IU_{1-\frac{\alpha}{2}}^{1-\frac{\tau}{2}}(x)$ in Equations (51) and (52).

$$P_{\mathcal{S}}\left(P_{Y(x)}(Y(x) \in IL_{\frac{\alpha}{2}}^{1-\frac{\tau}{2}}(x)|\mathcal{S}) \leq \frac{\alpha}{2}\right) = 1-\frac{\tau}{2}, \text{ where } IL_{\frac{\alpha}{2}}^{1-\frac{\tau}{2}}(x) =]-\infty, L_{\frac{\alpha}{2}}^{1-\frac{\tau}{2}}(x)], \quad (51)$$

$$P_{\mathcal{S}}\left(P_{Y(x)}(Y(x) \in IU_{1-\frac{\alpha}{2}}^{1-\frac{\tau}{2}}(x)|\mathcal{S}) \geq 1-\frac{\alpha}{2}\right) = 1-\frac{\tau}{2}, \text{ where } IU_{1-\frac{\alpha}{2}}^{1-\frac{\tau}{2}}(x) =]-\infty, U_{1-\frac{\alpha}{2}}^{1-\frac{\tau}{2}}(x)]. \quad (52)$$

In Equations (51) and (52), $L_{\frac{\alpha}{2}}^{1-\frac{\tau}{2}}(x)$ denotes an lower confidence bound on the $\frac{\alpha}{2}$-regression quantile at point x. This confidence bound must, a proportion $1-\frac{\tau}{2}$ of the time, cover the $\frac{\alpha}{2}$ quantile of $Y(x)$. $L_{1-\frac{\alpha}{2}}^{1-\frac{\tau}{2}}(x)$ denotes an upper confidence bound on the regression quantile at point x and it must, a proportion $1-\frac{\tau}{2}$ of the time, cover the $1-\frac{\alpha}{2}$ quantile of $Y(x)$.

Note that these confidence statements are made on two different models, and so we cannot use them directly to construct two-sided confidence intervals. However, by applying the Bonferroni inequality, one can merge the pair of $(1-\frac{\tau}{2})$ confidence intervals to obtain a joint confidence statement with a probability greater than or equal to $\gamma = 1-\tau$. Equation (53) describes this combination.

$$P_{\mathcal{S}}\left(P_{Y(x)}\left(Y(x) \in IQ_{1-\alpha}^{1-\tau}(x)\right) \geq 1-\alpha\right) \geq 1-\tau, \quad (53)$$
$$\text{where } IQ_{1-\alpha}^{1-\tau}(x) = [L_{\frac{\alpha}{2}}^{1-\frac{\tau}{2}}(x), U_{1-\frac{\alpha}{2}}^{1-\frac{\tau}{2}}(x)].$$

.4 Proof of Proposition 1

Proof: The mentioned assumptions lead to assume that the prediction error has a normal distribution and its variance is approximately the same in the neighborhood of x^*. Let x^* denote the query point and let $\varepsilon_{x^*}^{pred}$ denote its prediction error, then we have:

$$\varepsilon_{x^*}^{pred} = \varepsilon + f(x^*) - \hat{f}(x^*),$$

which results in:

$$\varepsilon_{x^*}^{pred} \sim \mathcal{N}(-Bias_{\hat{f}(x^*)}, \sigma_{x^*}^2 + \sigma_{\hat{f}(x^*)}^2). \quad (54)$$

The tolerance interval of the prediction error, is denoted by

$$I(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T} = [L(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T}, U(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T}] = ICentered(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T} - \hat{bias}_{\hat{f}(x^*)},$$

$$\widehat{bias}_{\hat{f}(x^*)} = \frac{1}{card(Kset_{x^*})}\sum_{x_i \in Kset_{x^*}} \varepsilon_i^{pred},$$

where $ICentered(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T}$, $card(Kset_{x^*})$ and $\widehat{bias}_{\hat{f}(x^*)}$ are respectively the zero-centered version of $I(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T}$, the cardinal of $Kset_{x^*}$ and the sample bias of $\hat{f}(x^*)$. Thus $I(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T}$ takes into account two kinds of uncertainties: the regression's method uncertainty and the observation error. Equation (54) shows that the prediction error $\varepsilon_{x^*}^{pred}$ has a normal distribution with the unknown mean $-Bias_{\hat{f}(x^*)}$. The prediction error tolerance interval $I(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T}$ is constructed based on the $Kset_{x^*}$, which is a finite sample size, so it is centered on the sample bias $-\widehat{bias}_{\hat{f}(x^*)}$. However, because of its definition, $I(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T}$ is guaranteed with confidence level γ, to contain at least a proportion β of the normal distribution of the prediction error at x^*. Hence we have:

$$P_{\mathcal{T}}\left(P_{\varepsilon}\left(L(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T} \leq \varepsilon_{x^*}^{pred} \leq U(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T} \middle| \mathcal{T}\right) \geq \beta\right) = \gamma,$$

where $\mathcal{T} = (\hat{f}(x^*), \sigma_{x^*})$ is the estimated vector at point x^*. This equation can be rewritten

$$\begin{aligned} &P_{\mathcal{T}}\left(P_{\varepsilon}\left(L(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T} \leq \varepsilon + f(x^*) - \hat{f}(x^*) \leq U(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T} \middle| \mathcal{T} \geq \beta\right)\right) \\ &= P_{\mathcal{T}}\left(P_{\varepsilon}\left(\hat{f}(x^*) + L(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T} \leq Y(x^*) \leq \hat{f}(x^*) + U(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T} \middle| \mathcal{T}\right) \geq \beta\right) \\ &= P_{\mathcal{T}}\left(P_{\varepsilon}\left(Y(x^*) \in (\hat{f}(x^*) + I(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T}) \middle| \mathcal{T}\right) \geq \beta\right) = \gamma. \end{aligned} \tag{55}$$

Equation (55) means that, by taking our assumptions, the tolerance interval for the response variable is computed by adding the local linear regression estimate to the tolerance interval on the prediction error:

$$I(x^*)_{\gamma,\beta}^{T} = \hat{f}(x^*) + I(\varepsilon_{x^*}^{pred})_{\gamma,\beta}^{T} \blacksquare$$

Even though we have a biased prediction, our tolerance interval contains the desired proportion of the conditional distribution of the response variable. This is because the tolerance intervals are computed on the prediction error and the prediction error is centered on negative bias. Proofs of proposition